

(12) **United States Patent**
Cole et al.

(10) **Patent No.:** **US 10,805,593 B2**
(45) **Date of Patent:** **Oct. 13, 2020**

(54) **METHODS AND APPARATUS FOR RECEIVING AND/OR USING REDUCED RESOLUTION IMAGES**

(71) Applicant: **Nevermind Capital LLC**, Wilmington, DE (US)

(72) Inventors: **David Cole**, Aliso Viejo, CA (US); **Alan McKay Moss**, Laguna Beach, CA (US); **Hector M Medina**, Laguna Beach, CA (US)

(73) Assignee: **Nevermind Capital LLC**, Wilmington, DE (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **16/708,420**

(22) Filed: **Dec. 9, 2019**

(65) **Prior Publication Data**
US 2020/0186779 A1 Jun. 11, 2020

Related U.S. Application Data
(60) Continuation of application No. 16/035,459, filed on Jul. 13, 2018, now Pat. No. 10,506,215, which is a (Continued)

(51) **Int. Cl.**
H04N 19/172 (2014.01)
H04N 13/189 (2018.01)
(Continued)

(52) **U.S. Cl.**
CPC **H04N 13/189** (2018.05); **G06T 15/04** (2013.01); **G06T 17/20** (2013.01);
(Continued)


(58) **Field of Classification Search**
None
See application file for complete search history.

(56) **References Cited**
U.S. PATENT DOCUMENTS
10,279,925 B2 5/2019 Cole et al.
10,432,910 B2 10/2019 Cole et al.
(Continued)

OTHER PUBLICATIONS
D'Angio, "Adaptive and Passive Non-Visual Driver Assistance Technologies for the Blind Driver Challenge" (Year: 2012).
Primary Examiner — Eileen M Adams
(74) *Attorney, Agent, or Firm* — Blank Rome LLP

(57) **ABSTRACT**
Methods and apparatus for using selective resolution reduction on images to be transmitted and/or used by a playback device are described. Prior to transmission one or more images of an environment are captured. Based on image content, motion detection and/or user input a resolution reduction operation is selected and performed. The reduced resolution image is communicated to a playback device along with information indicating a UV map corresponding to the selected resolution allocation that should be used by the playback device for rendering the communicated image. By changing the resolution allocation used and which UV map is used by the playback device different resolution allocations can be made with respect to different portions of the environment while allowing the number of pixels in transmitted images to remain constant. The playback device renders the individual images with the UV map corresponding to the resolution allocation used to generate the individual images.

14 Claims, 32 Drawing Sheets



Related U.S. Application Data

continuation of application No. 15/672,236, filed on Aug. 8, 2017, now Pat. No. 10,027,944, which is a division of application No. 15/046,327, filed on Feb. 17, 2016, now Pat. No. 9,729,850.

- (60) Provisional application No. 62/296,065, filed on Feb. 16, 2016, provisional application No. 62/262,374, filed on Dec. 2, 2015, provisional application No. 62/117,427, filed on Feb. 17, 2015.

(51) **Int. Cl.**

H04N 21/437 (2011.01)
H04N 21/2387 (2011.01)
H04N 21/2343 (2011.01)
H04N 5/232 (2006.01)
H04N 19/136 (2014.01)
H04N 19/597 (2014.01)
H04N 13/279 (2018.01)
H04N 13/275 (2018.01)
H04N 13/243 (2018.01)
H04N 13/194 (2018.01)
H04N 13/161 (2018.01)
H04N 13/139 (2018.01)
H04N 13/106 (2018.01)
H04N 21/81 (2011.01)
G06T 15/04 (2011.01)
H04N 21/426 (2011.01)
H04N 21/435 (2011.01)
H04N 21/44 (2011.01)
G06T 17/20 (2006.01)
H04N 19/44 (2014.01)
H04N 21/472 (2011.01)
H04N 13/00 (2018.01)

(52) **U.S. Cl.**

CPC *H04N 5/23238* (2013.01); *H04N 13/106* (2018.05); *H04N 13/139* (2018.05); *H04N 13/161* (2018.05); *H04N 13/194* (2018.05); *H04N 13/243* (2018.05); *H04N 13/275* (2018.05); *H04N 13/279* (2018.05); *H04N 19/136* (2014.11); *H04N 19/172* (2014.11); *H04N 19/44* (2014.11); *H04N 19/597* (2014.11); *H04N 21/2387* (2013.01); *H04N 21/234363* (2013.01); *H04N 21/42653* (2013.01); *H04N 21/435* (2013.01); *H04N 21/437* (2013.01); *H04N 21/44016* (2013.01); *H04N 21/472* (2013.01); *H04N 21/816* (2013.01); *H04N 21/8146* (2013.01); *G06T 2215/12* (2013.01); *H04N 2013/0088* (2013.01)

(56)

References Cited

U.S. PATENT DOCUMENTS

10,438,369 B2	10/2019	Cole et al.	
10,477,186 B2	11/2019	Sheridan	
10,523,920 B2	12/2019	Cole et al.	
10,531,071 B2	1/2020	Cole et al.	
10,567,733 B2	2/2020	Cole et al.	
2009/0021522 A1*	1/2009	Burley	G06T 15/04 345/582
2013/0125155 A1	5/2013	Bhagavathy et al.	
2013/0278631 A1	10/2013	Border et al.	
2013/0300740 A1*	11/2013	Snyder	G06F 3/016 345/420
2015/0082181 A1	3/2015	Ames et al.	
2015/0264299 A1	9/2015	Leech et al.	

* cited by examiner







FIGURE 1



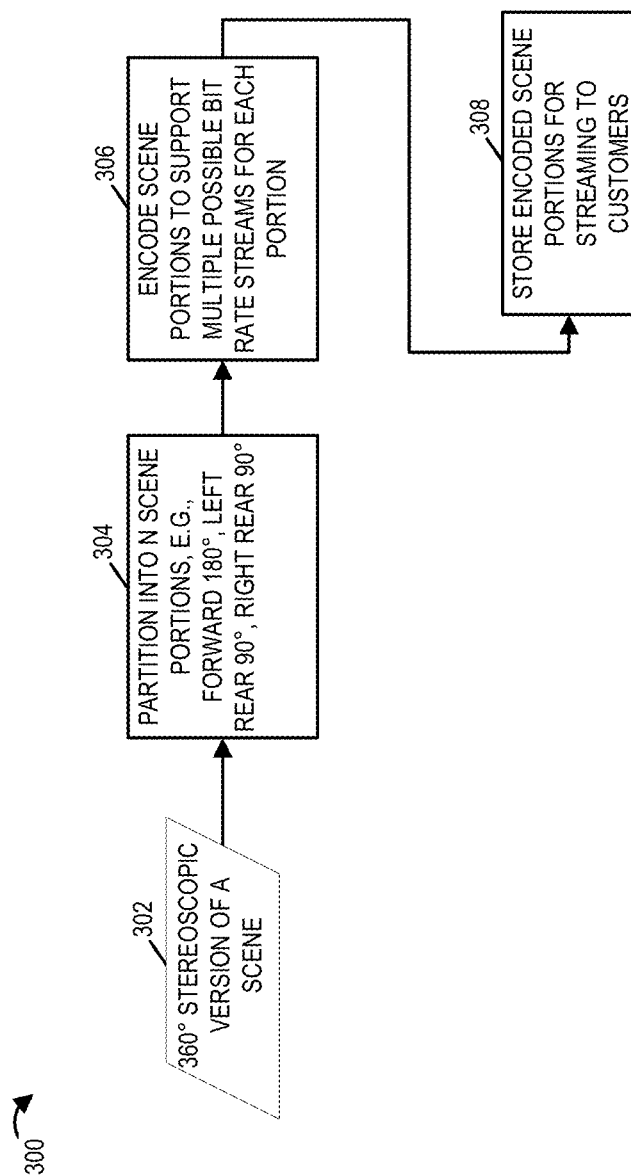


FIGURE 3




FIGURE 4

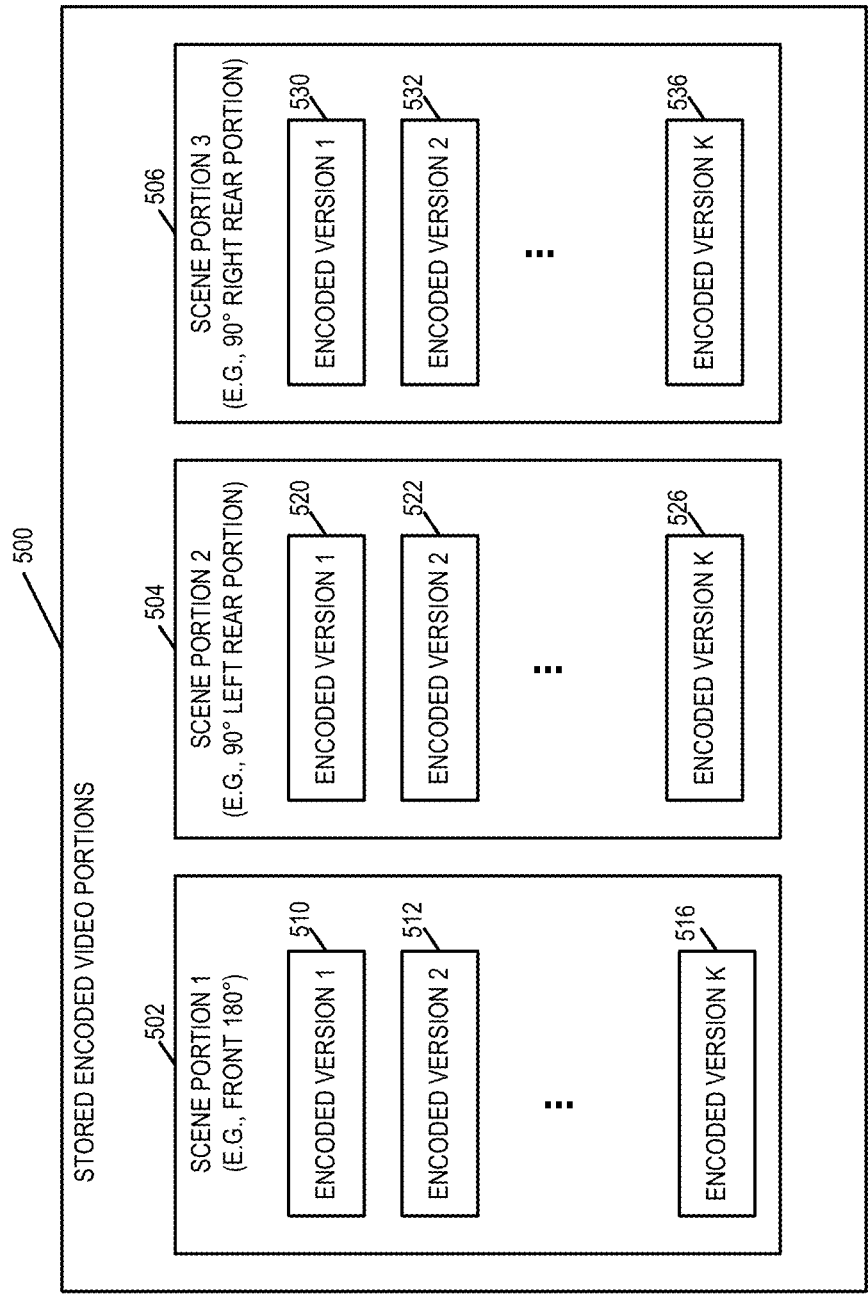


FIGURE 5

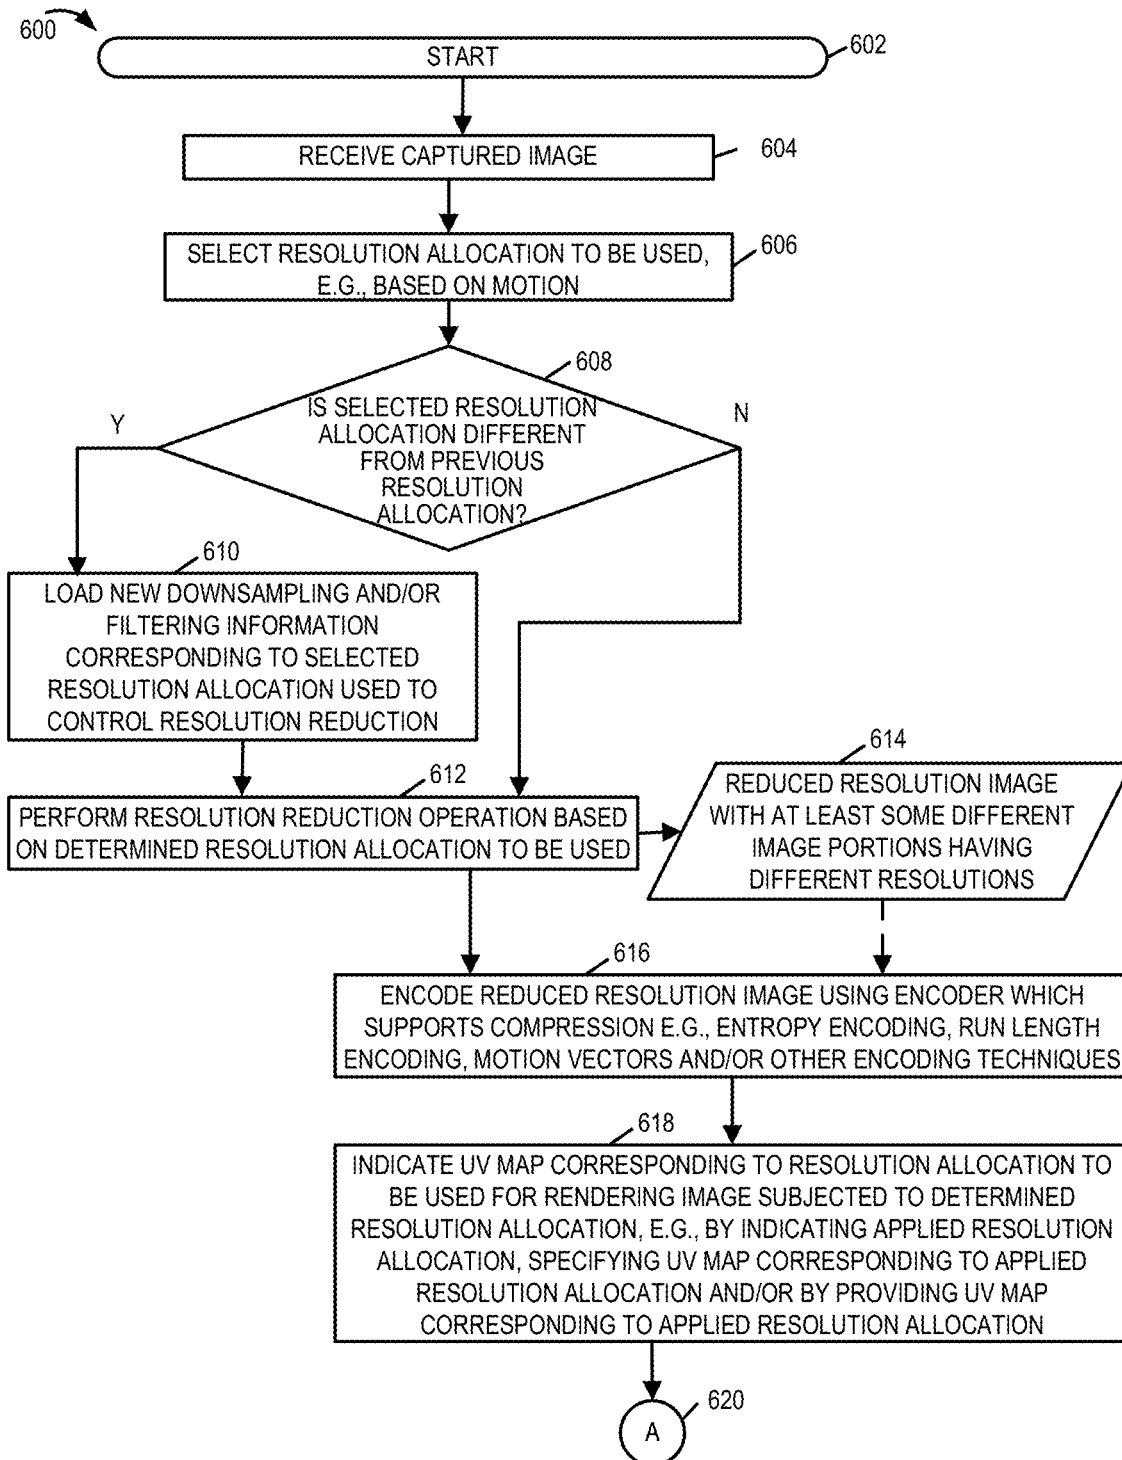


FIGURE 6A

FIGURE 6A

FIGURE 6B

FIGURE 6

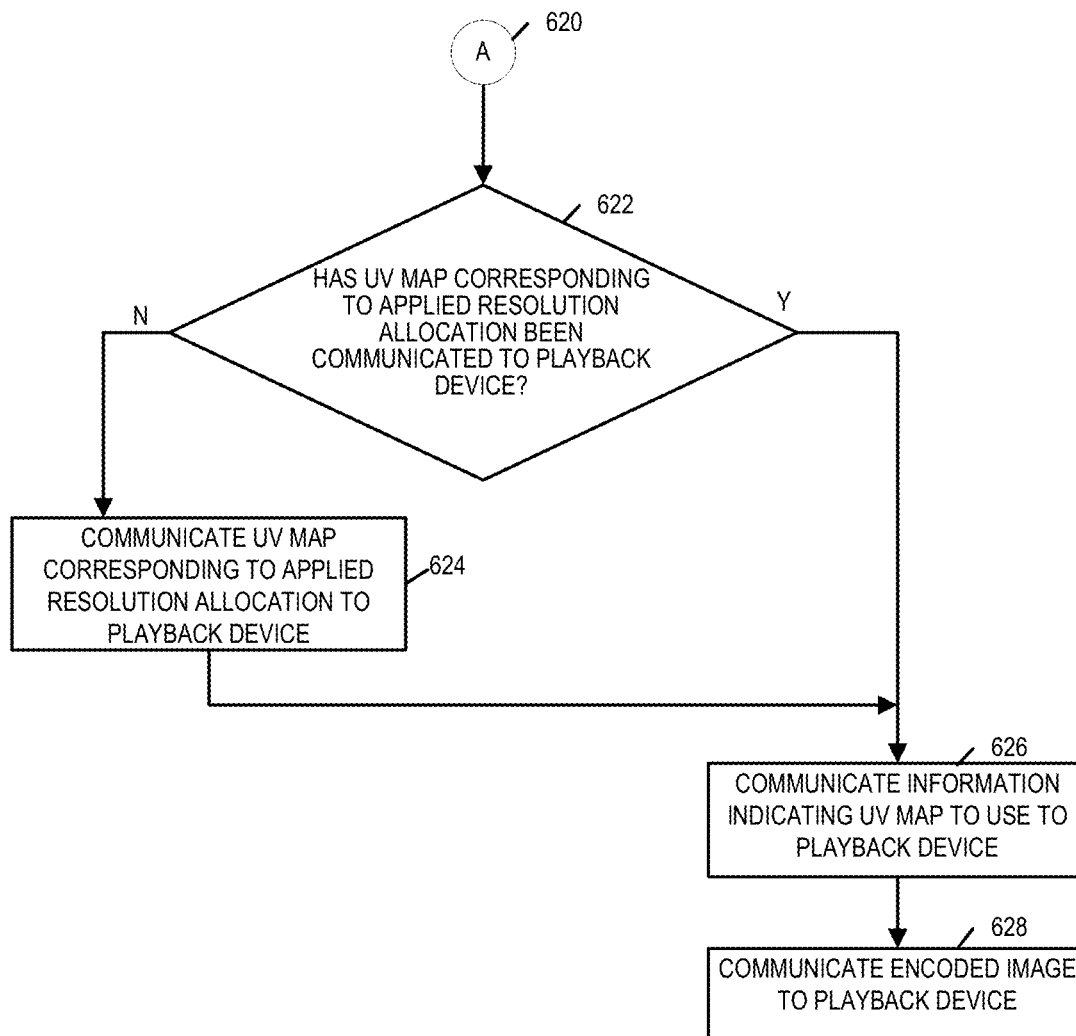


FIGURE 6B

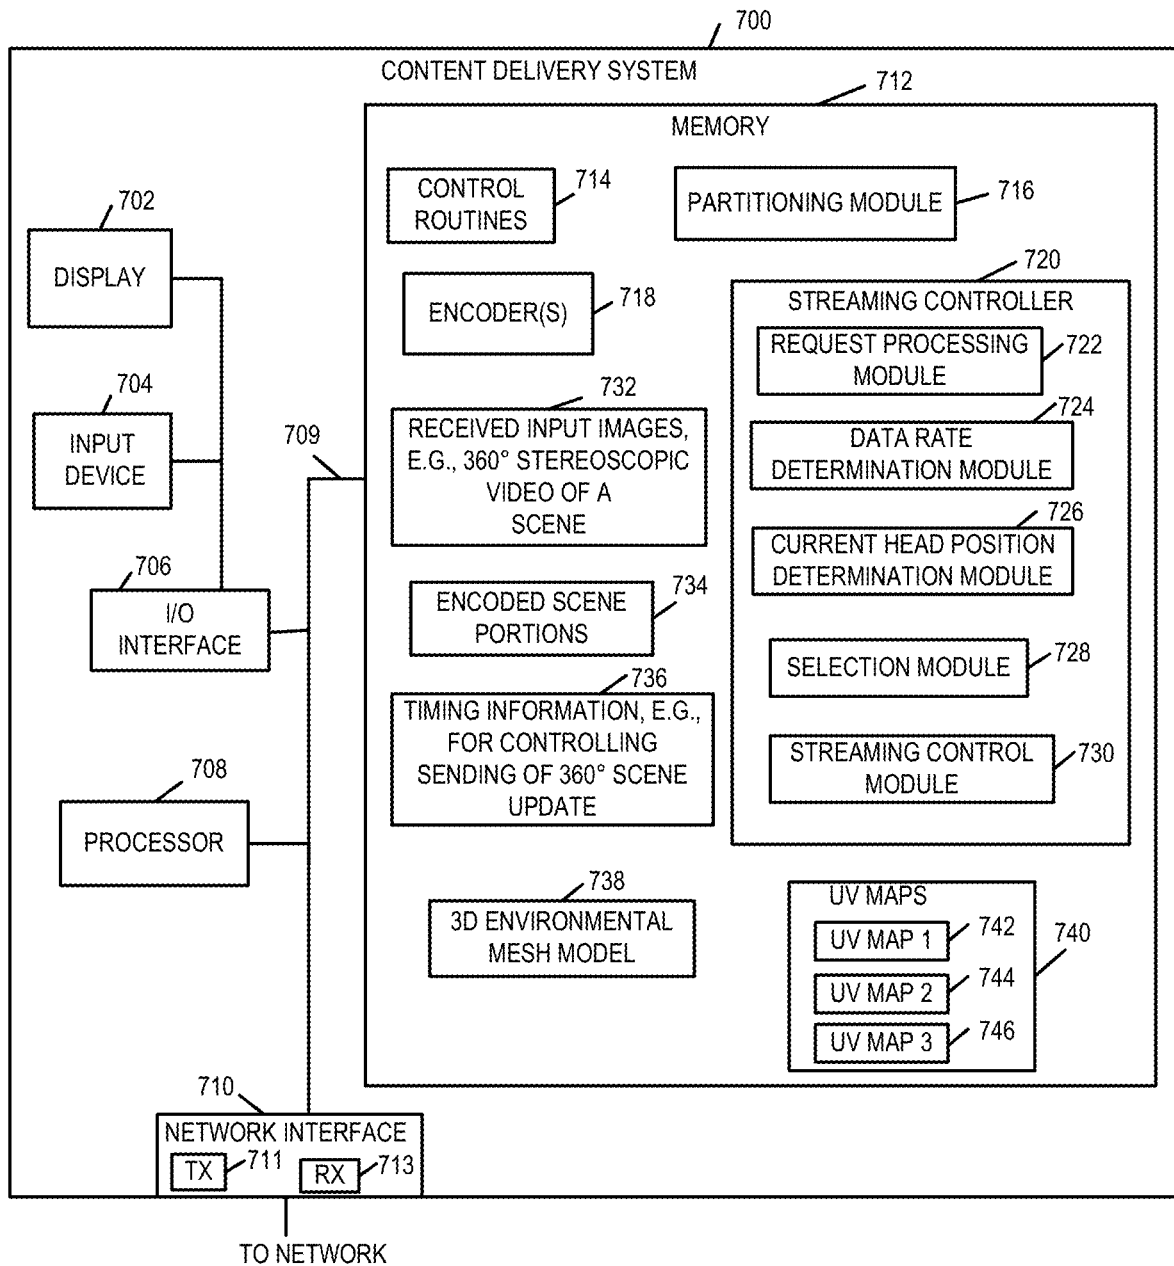


FIGURE 7

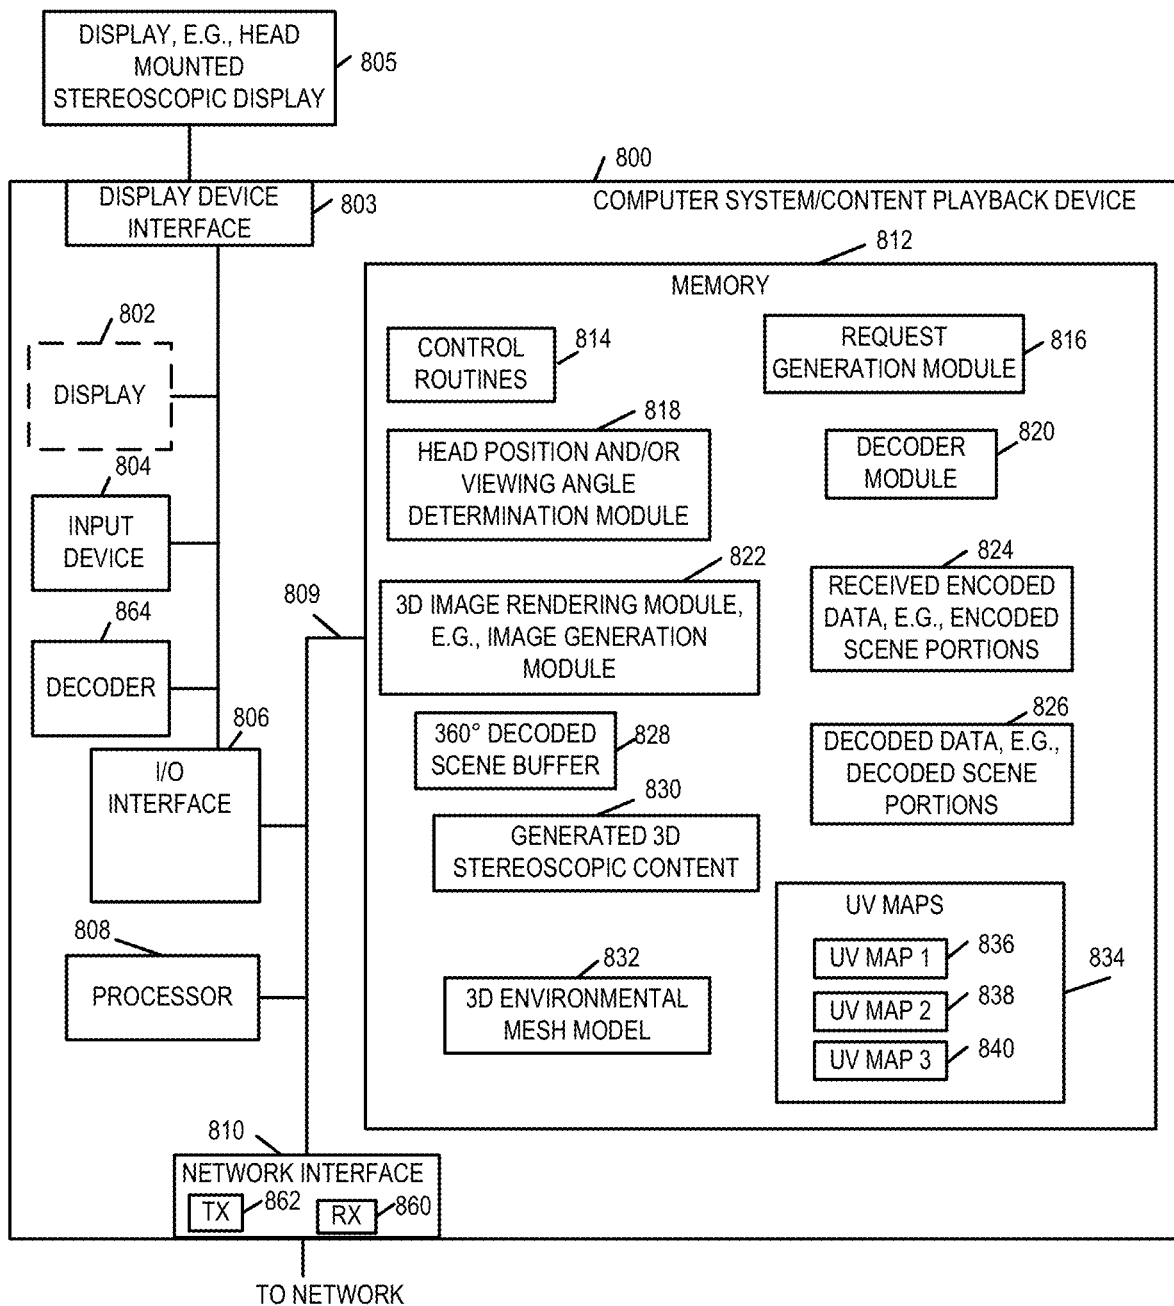


FIGURE 8

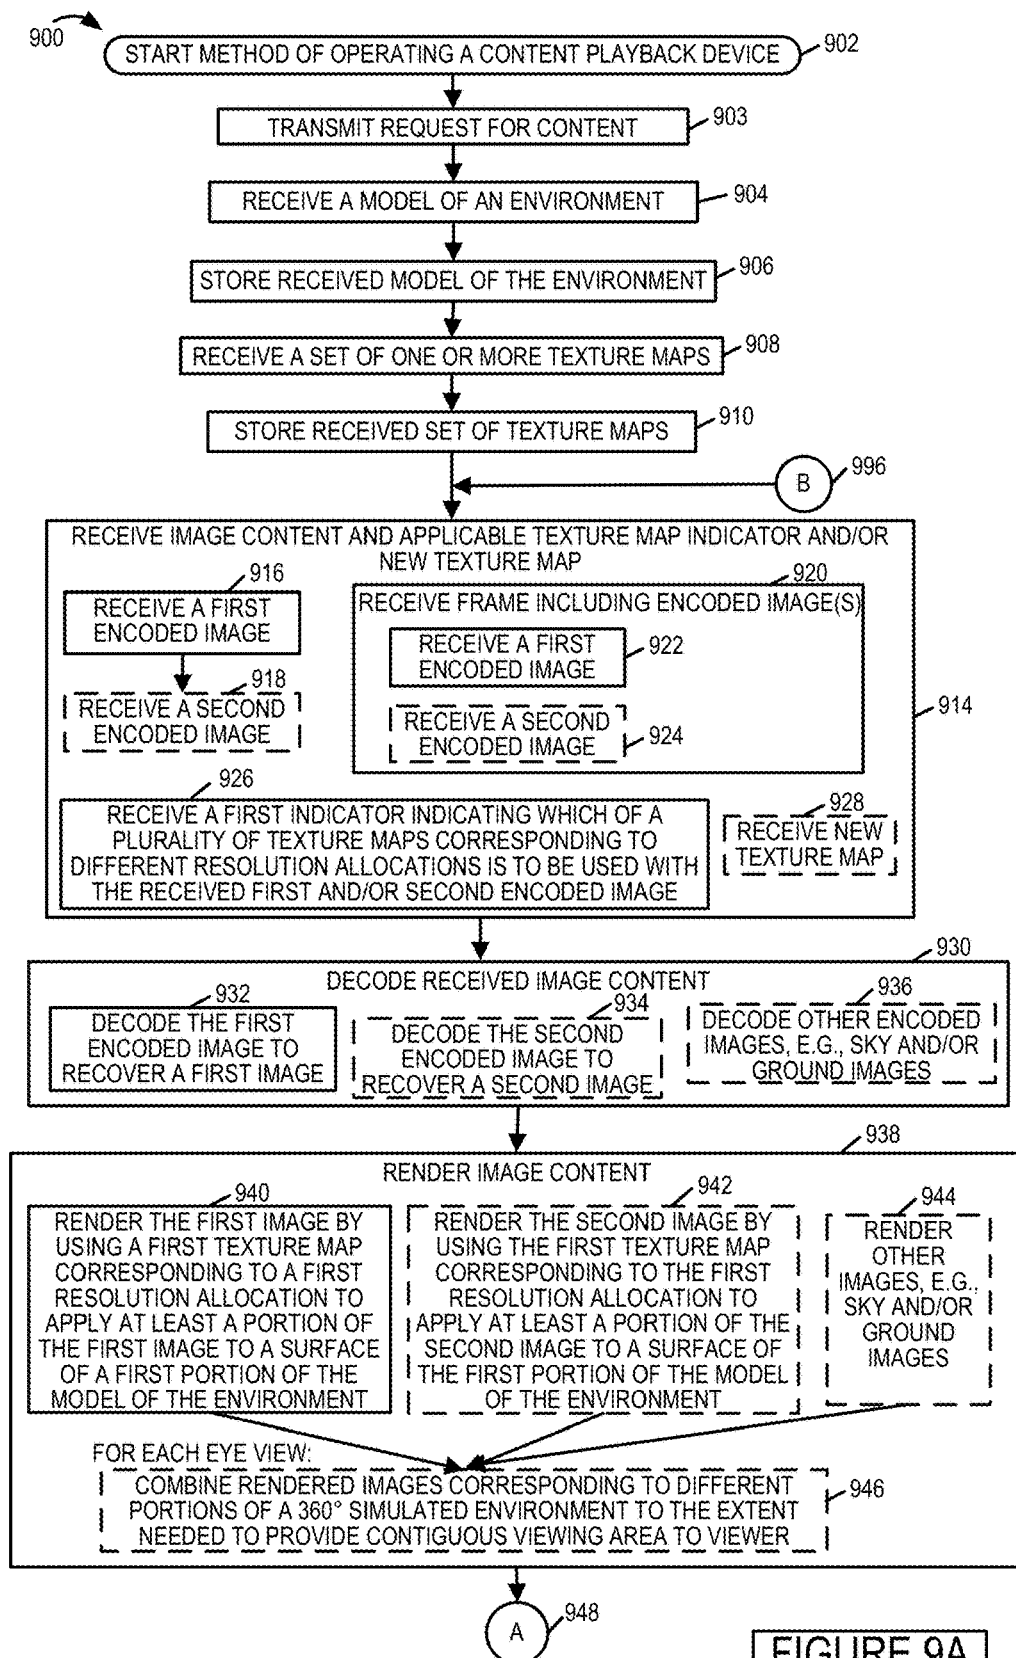


FIGURE 9A

FIGURE 9A
FIGURE 9B
FIGURE 9

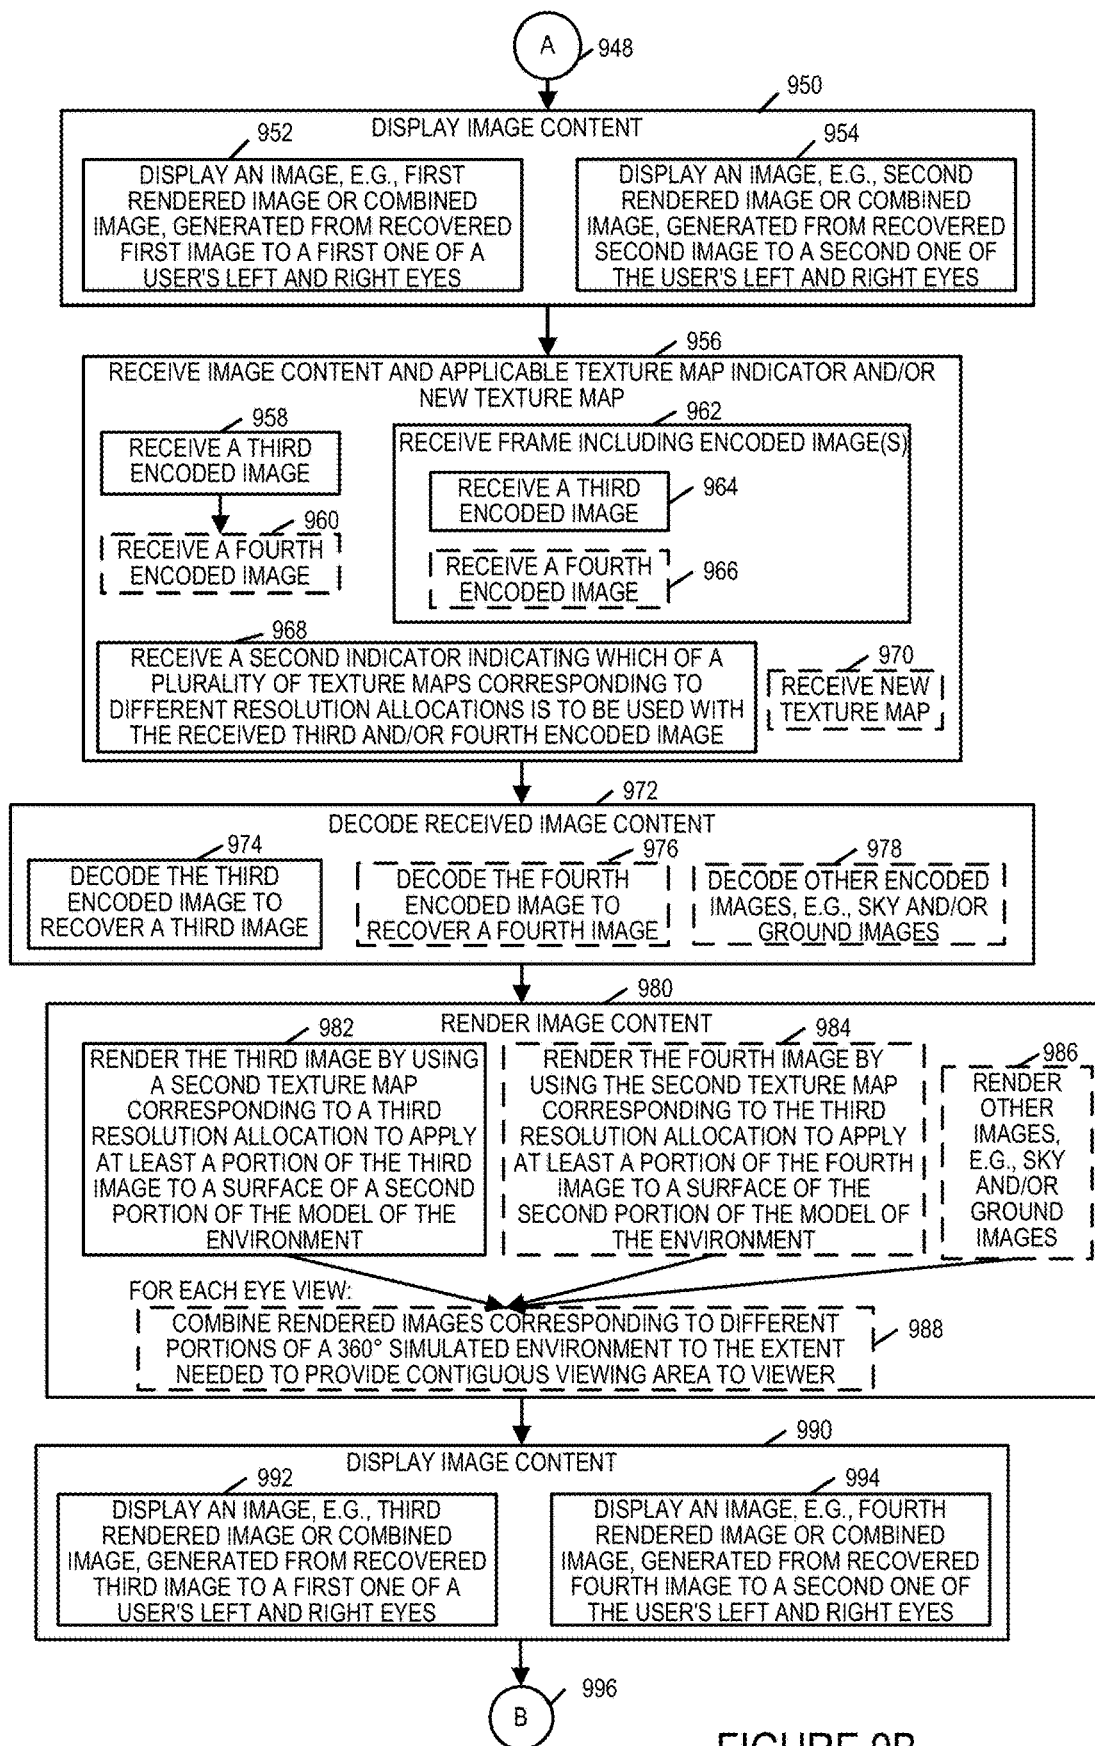


FIGURE 9B

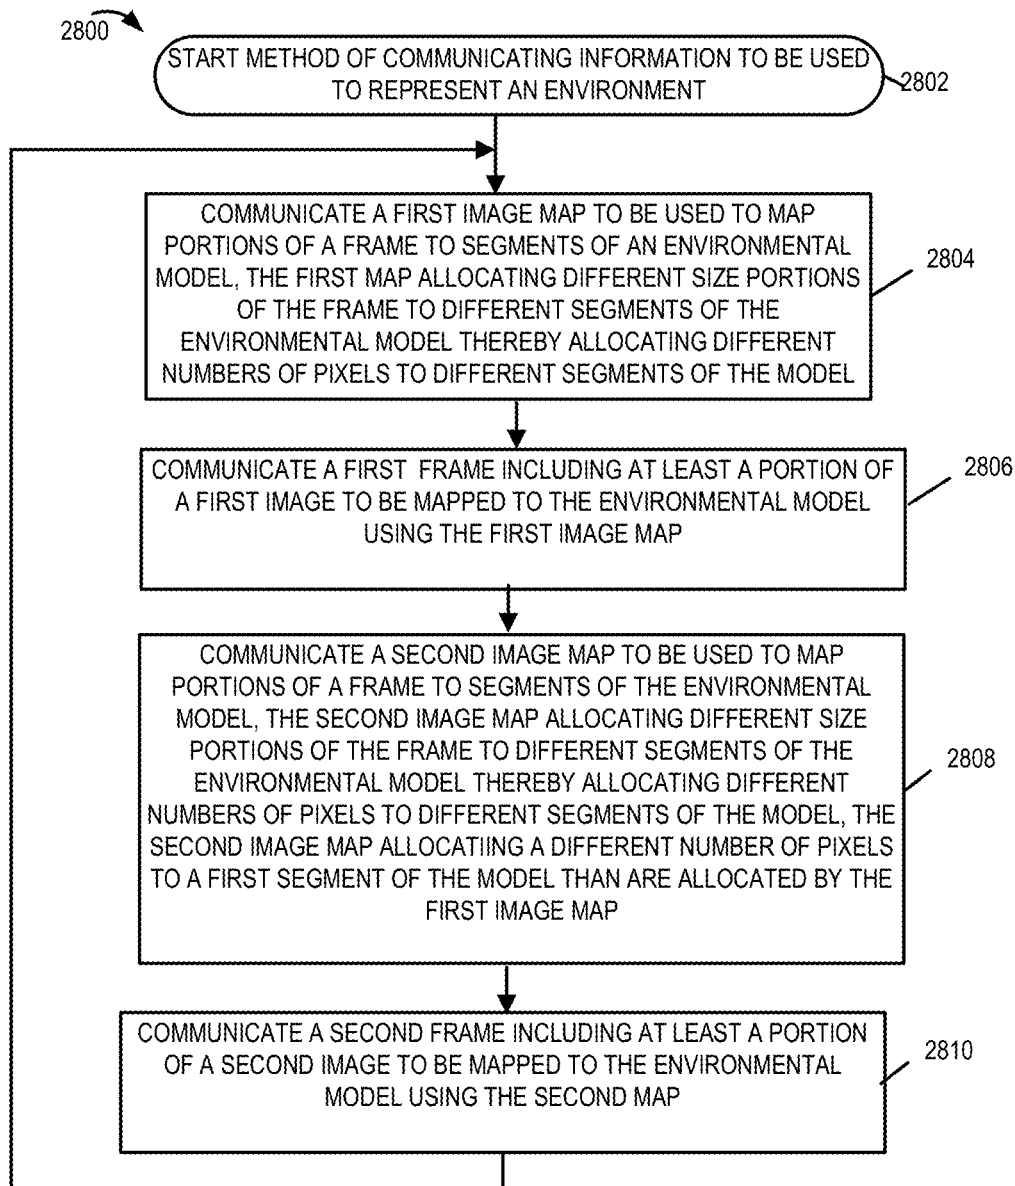


FIGURE 10

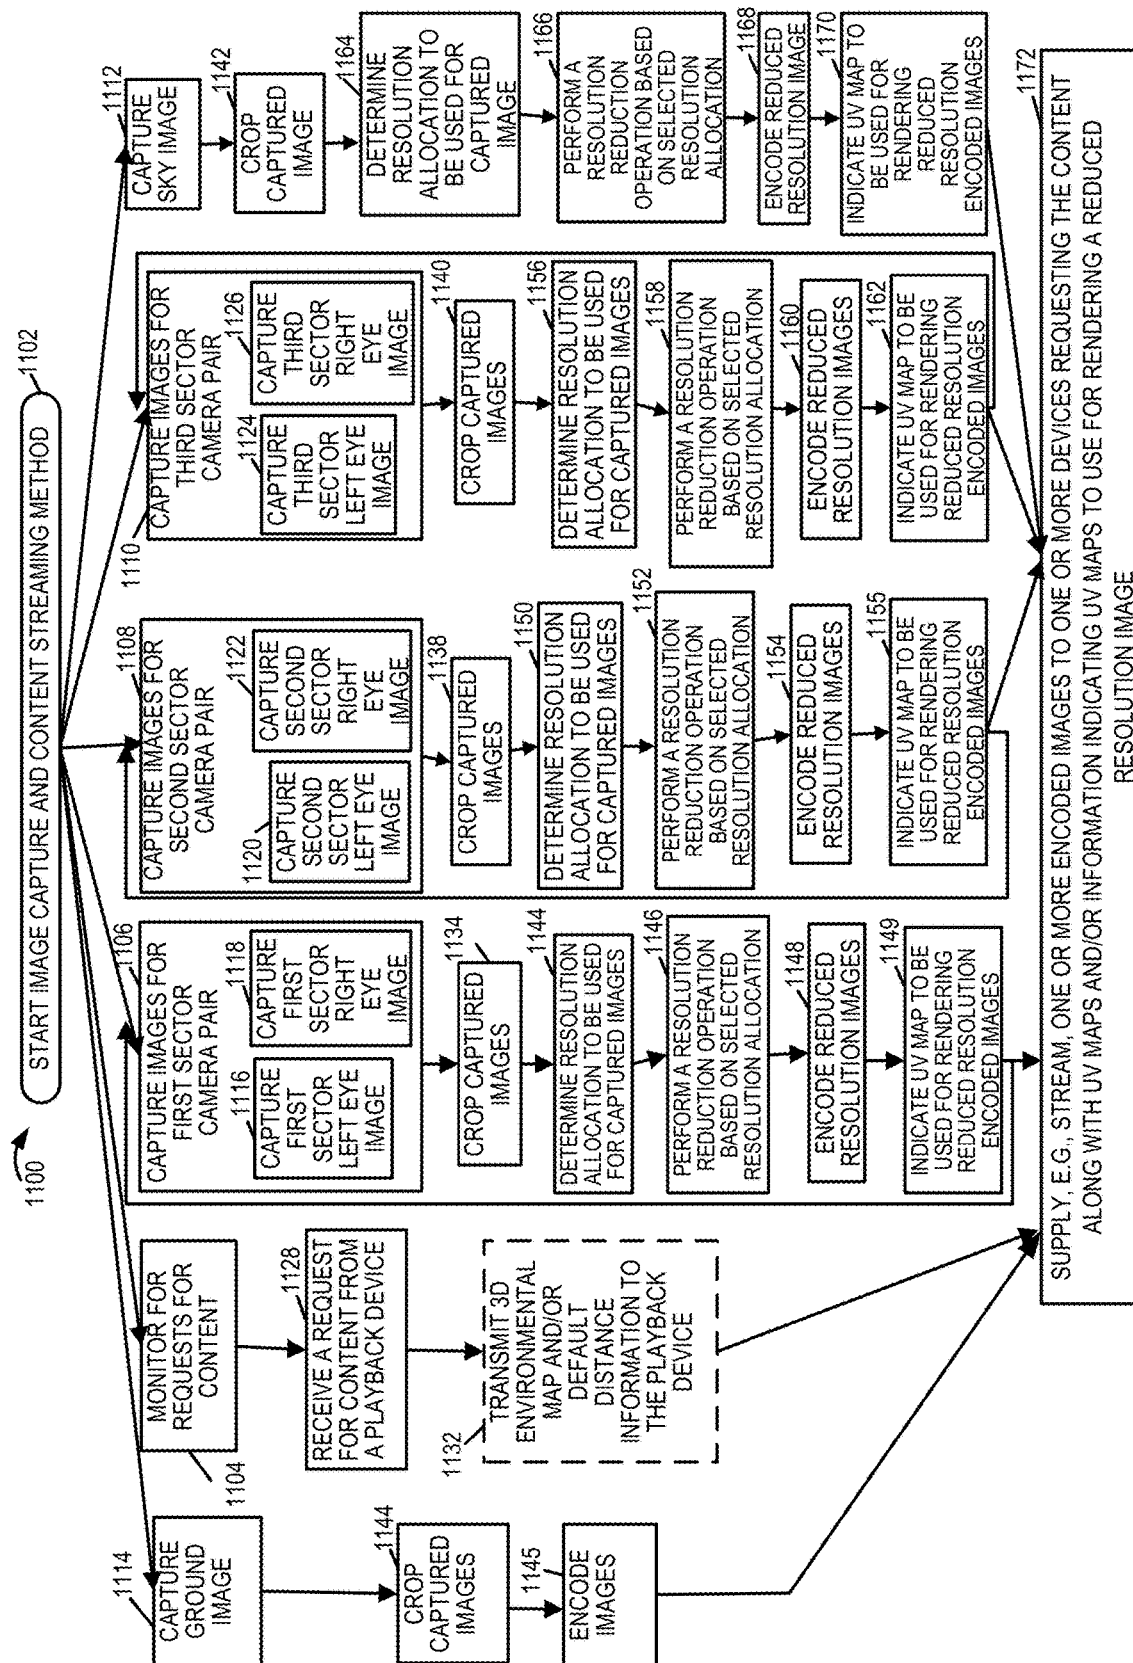


FIGURE 11

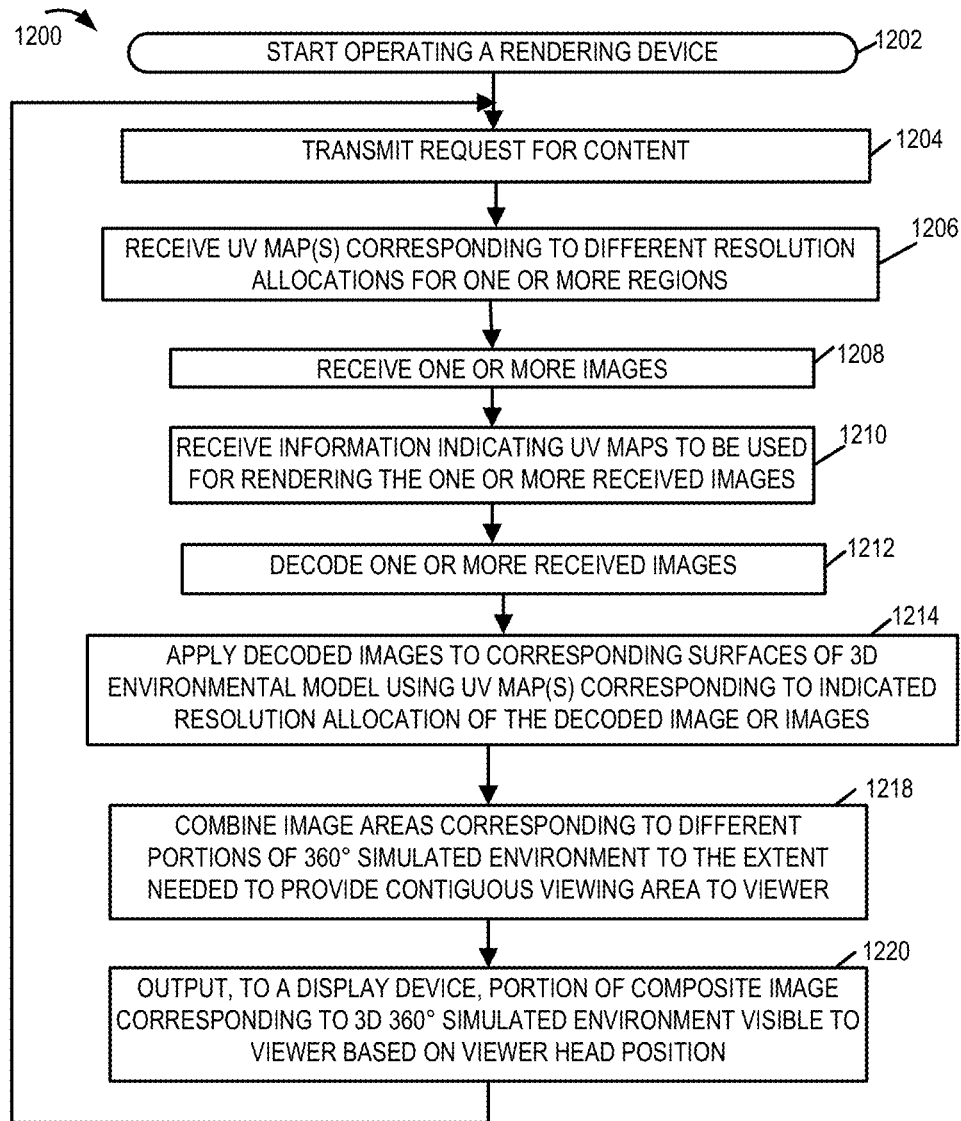


FIGURE 12




FIGURE 13

1400




FIGURE 14




FIGURE 15




FIGURE 16




FIGURE 17A




FIGURE 17B

2000





FIGURE 18



0-View
3D Model Mesh

FIGURE 19




FIGURE 20

2300 Selective Resolution Allocation in a Panoramic Image Map




FIGURE 21




FIGURE 22




FIGURE 23




FIGURE 24

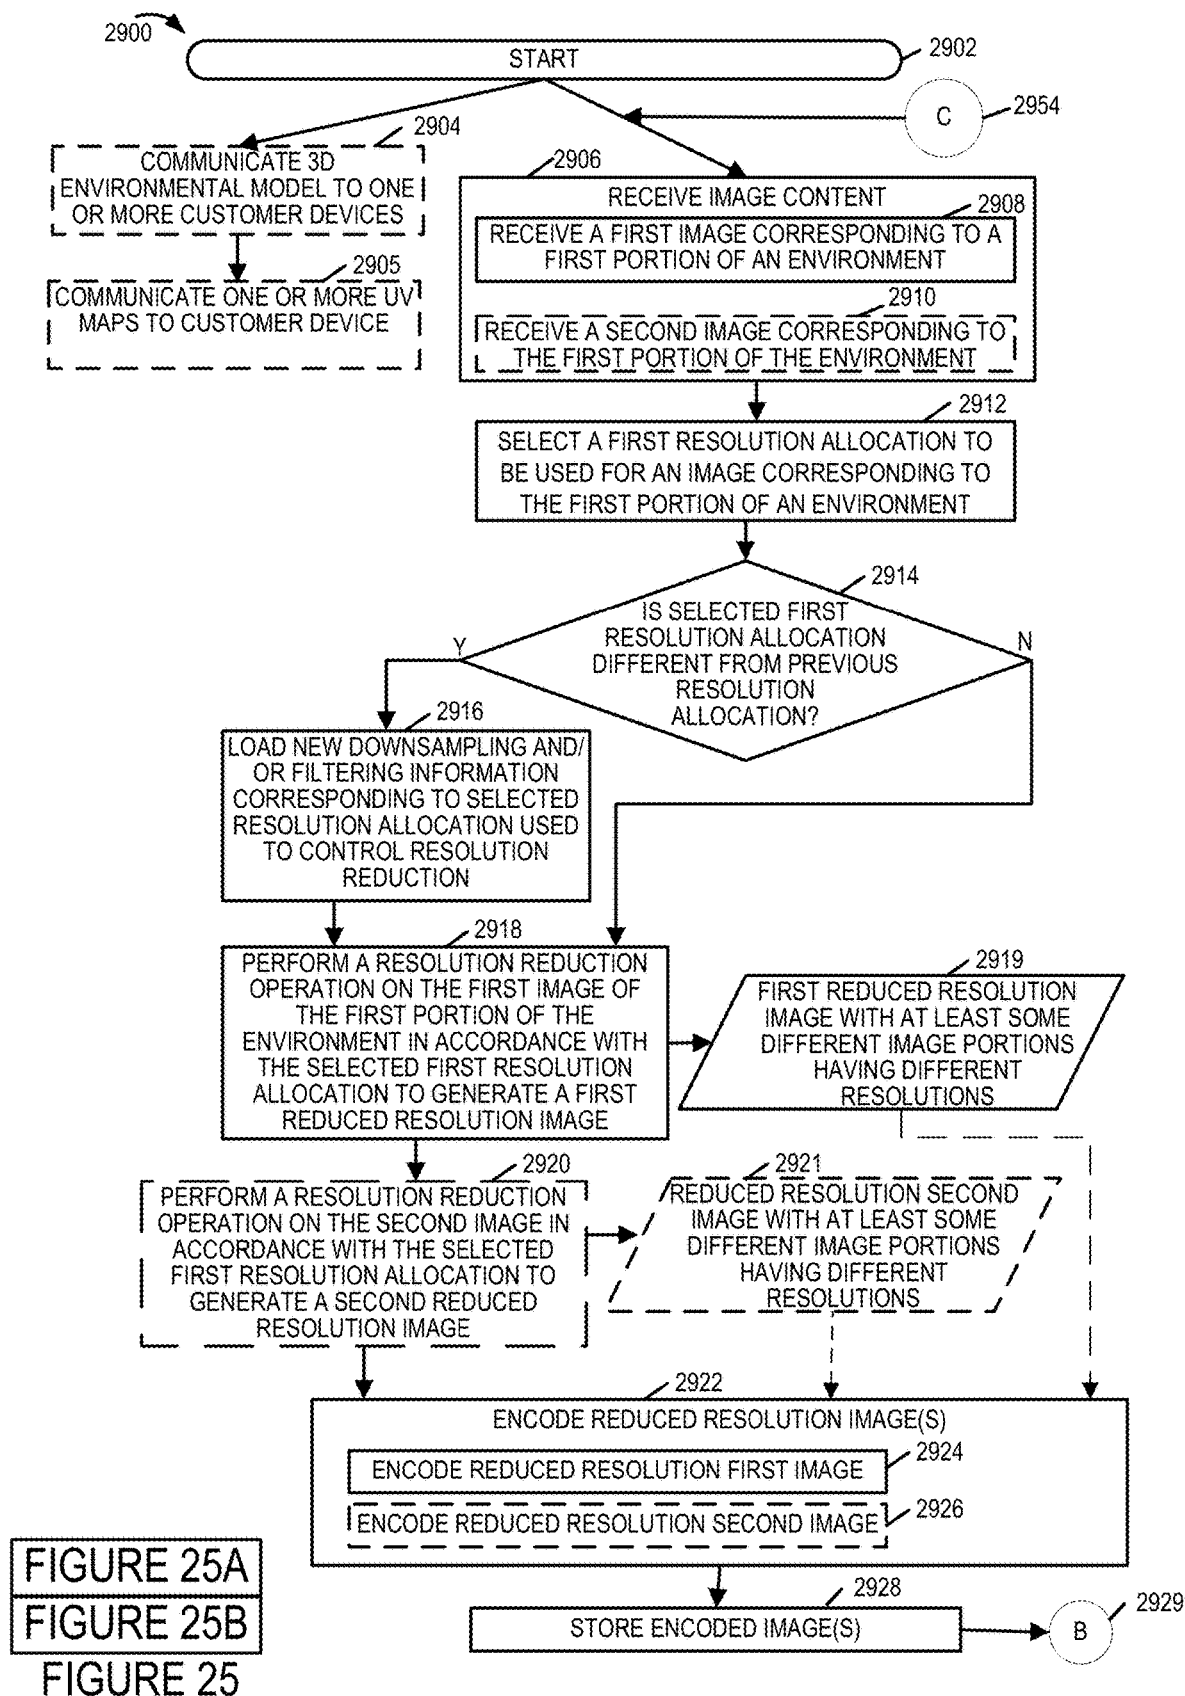


FIGURE 25A

FIGURE 25B

FIGURE 25

FIGURE 25A

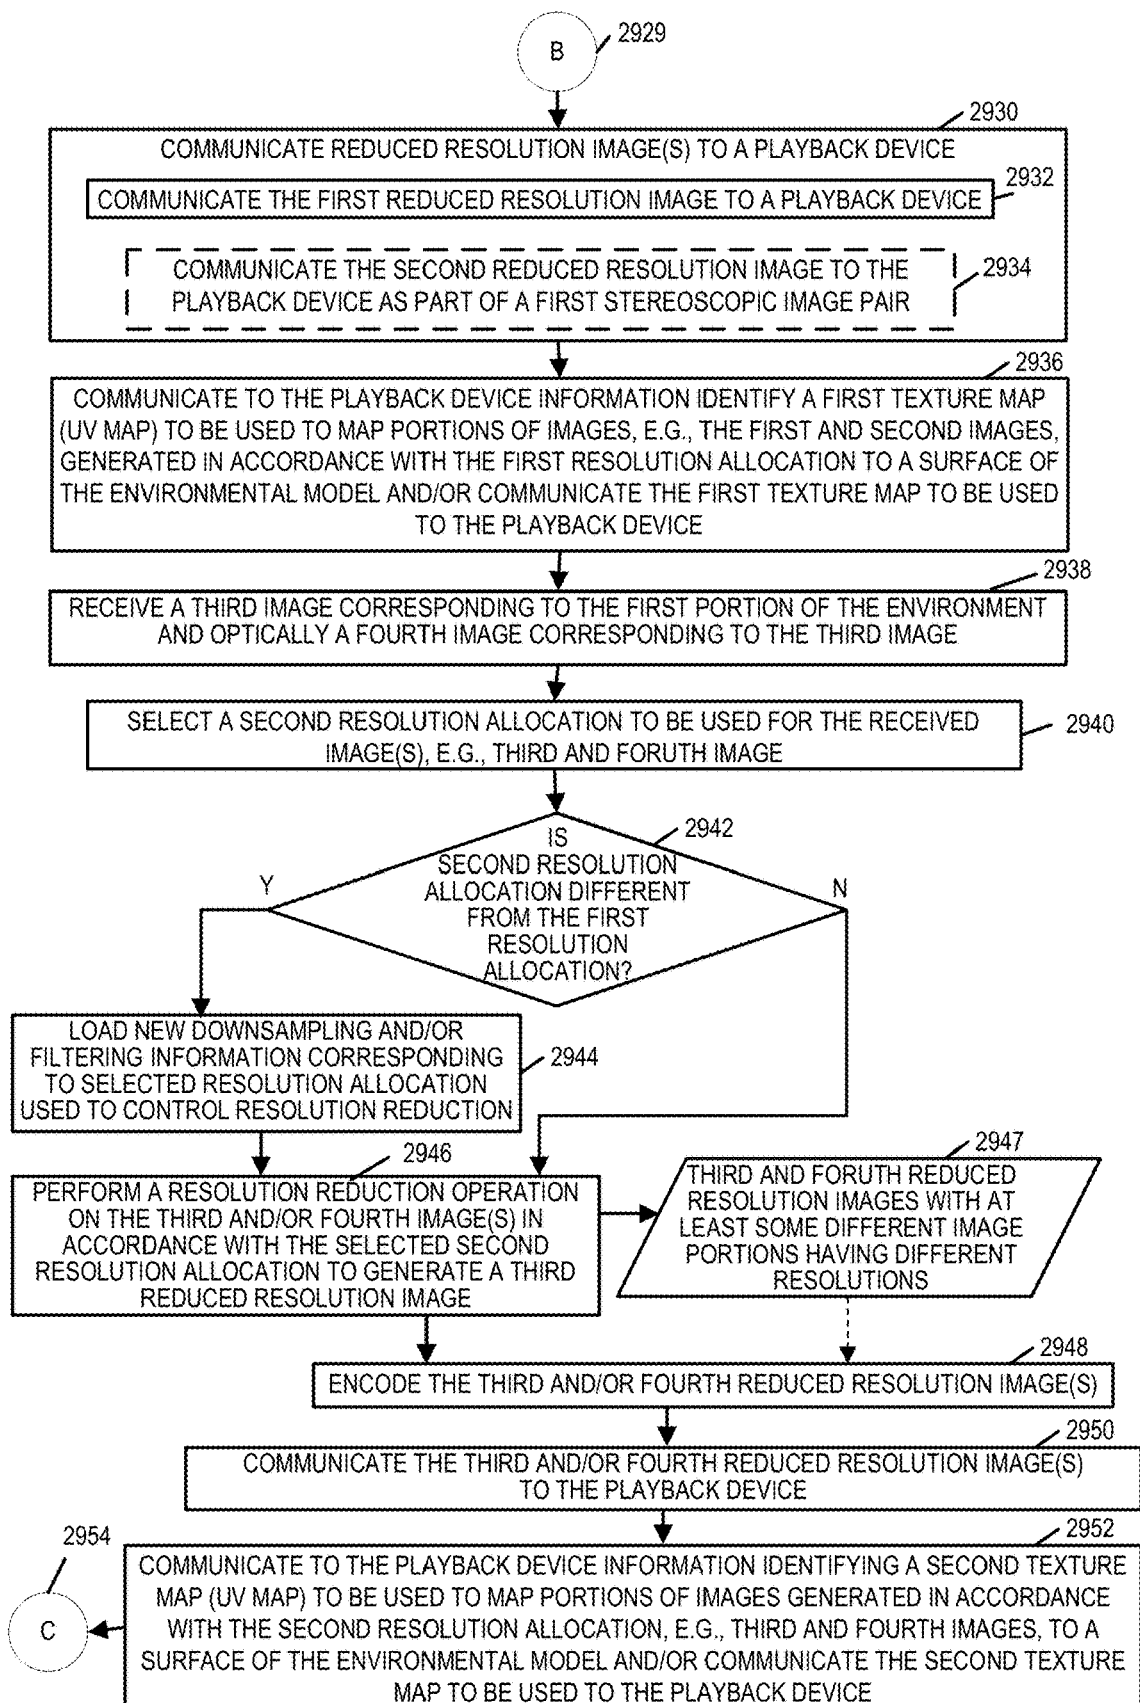


FIGURE 25B

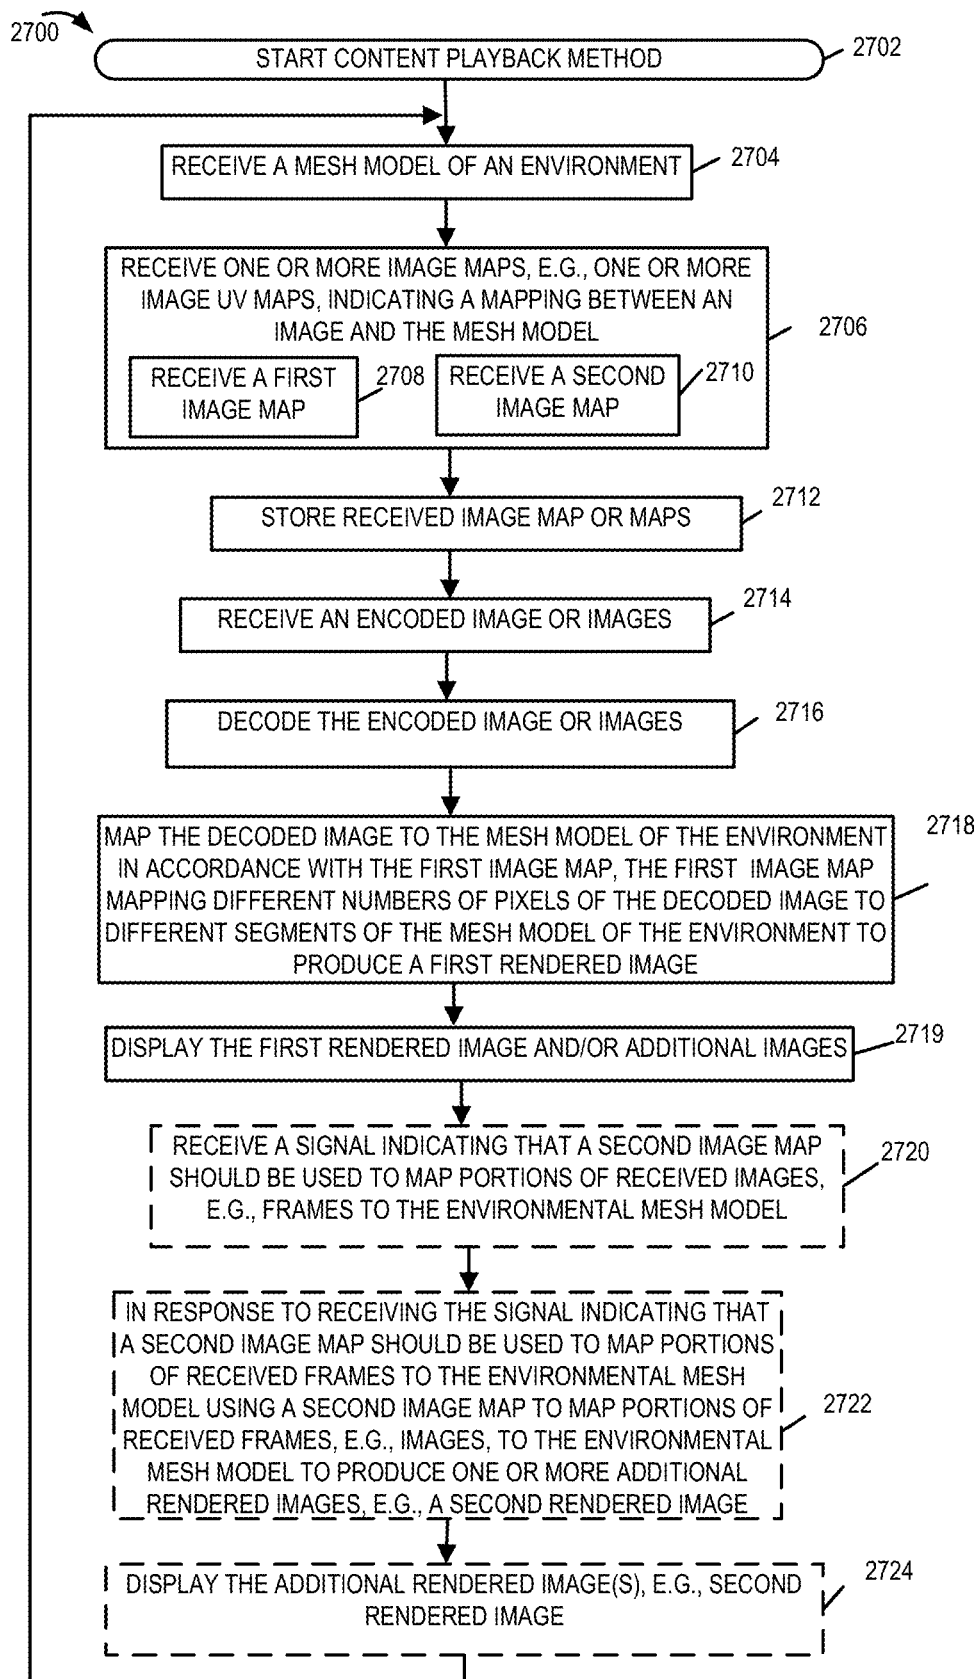


FIGURE 26

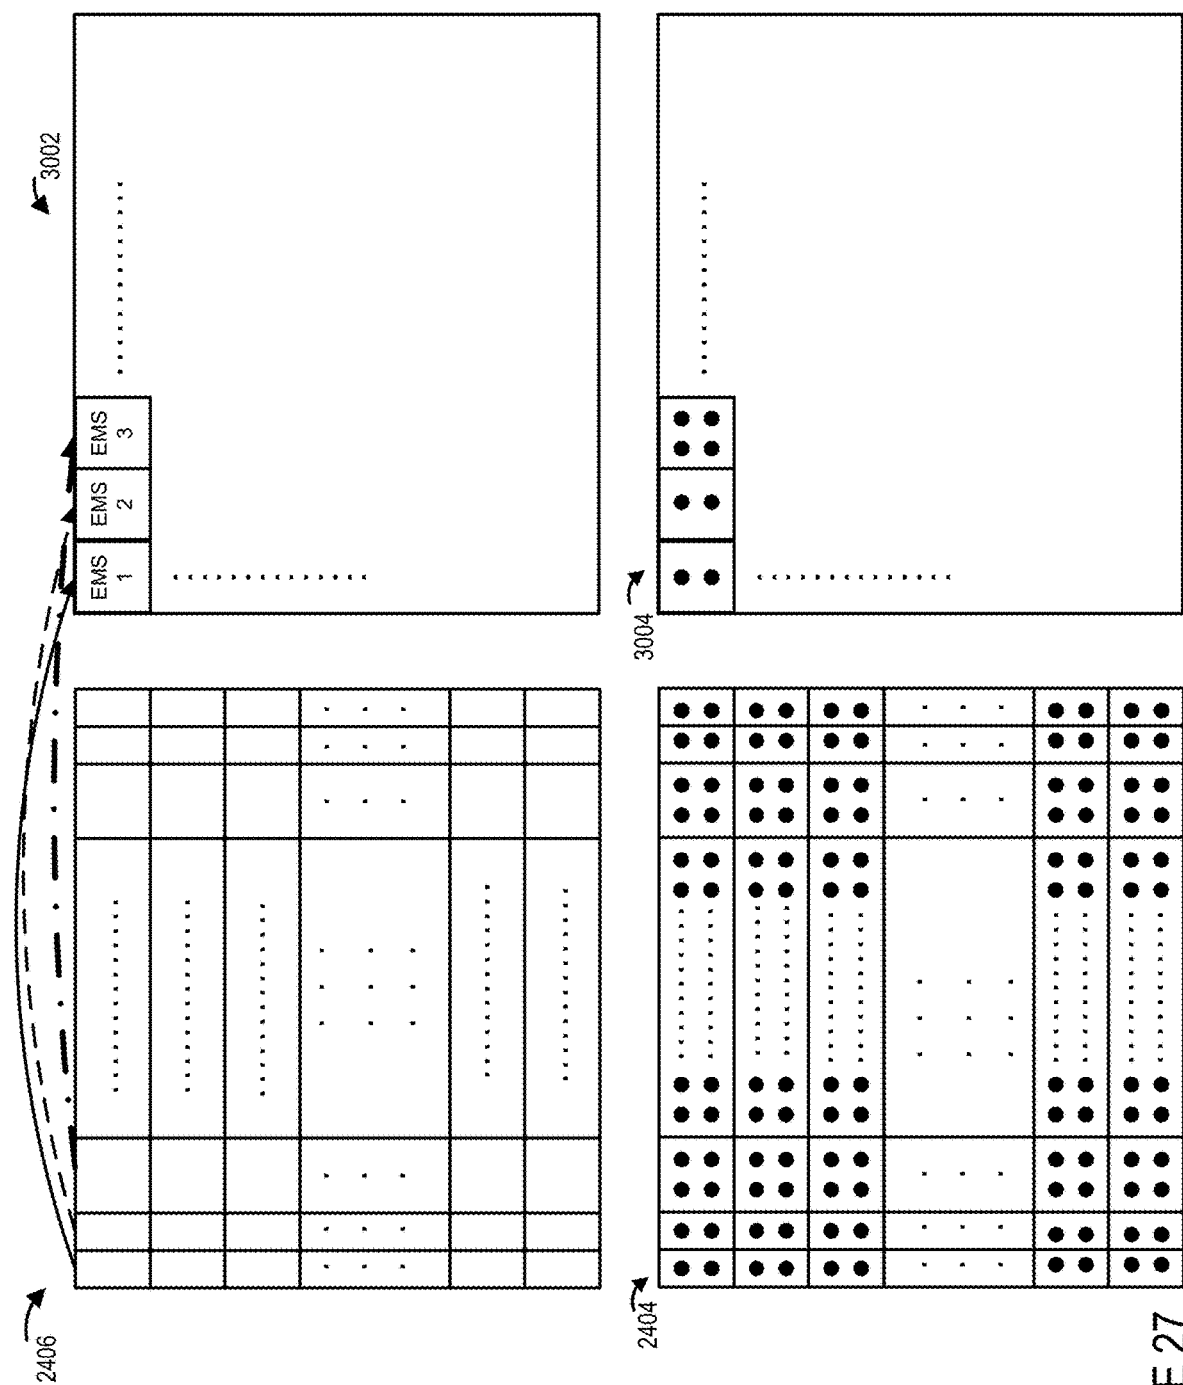


FIGURE 27




FIGURE 28

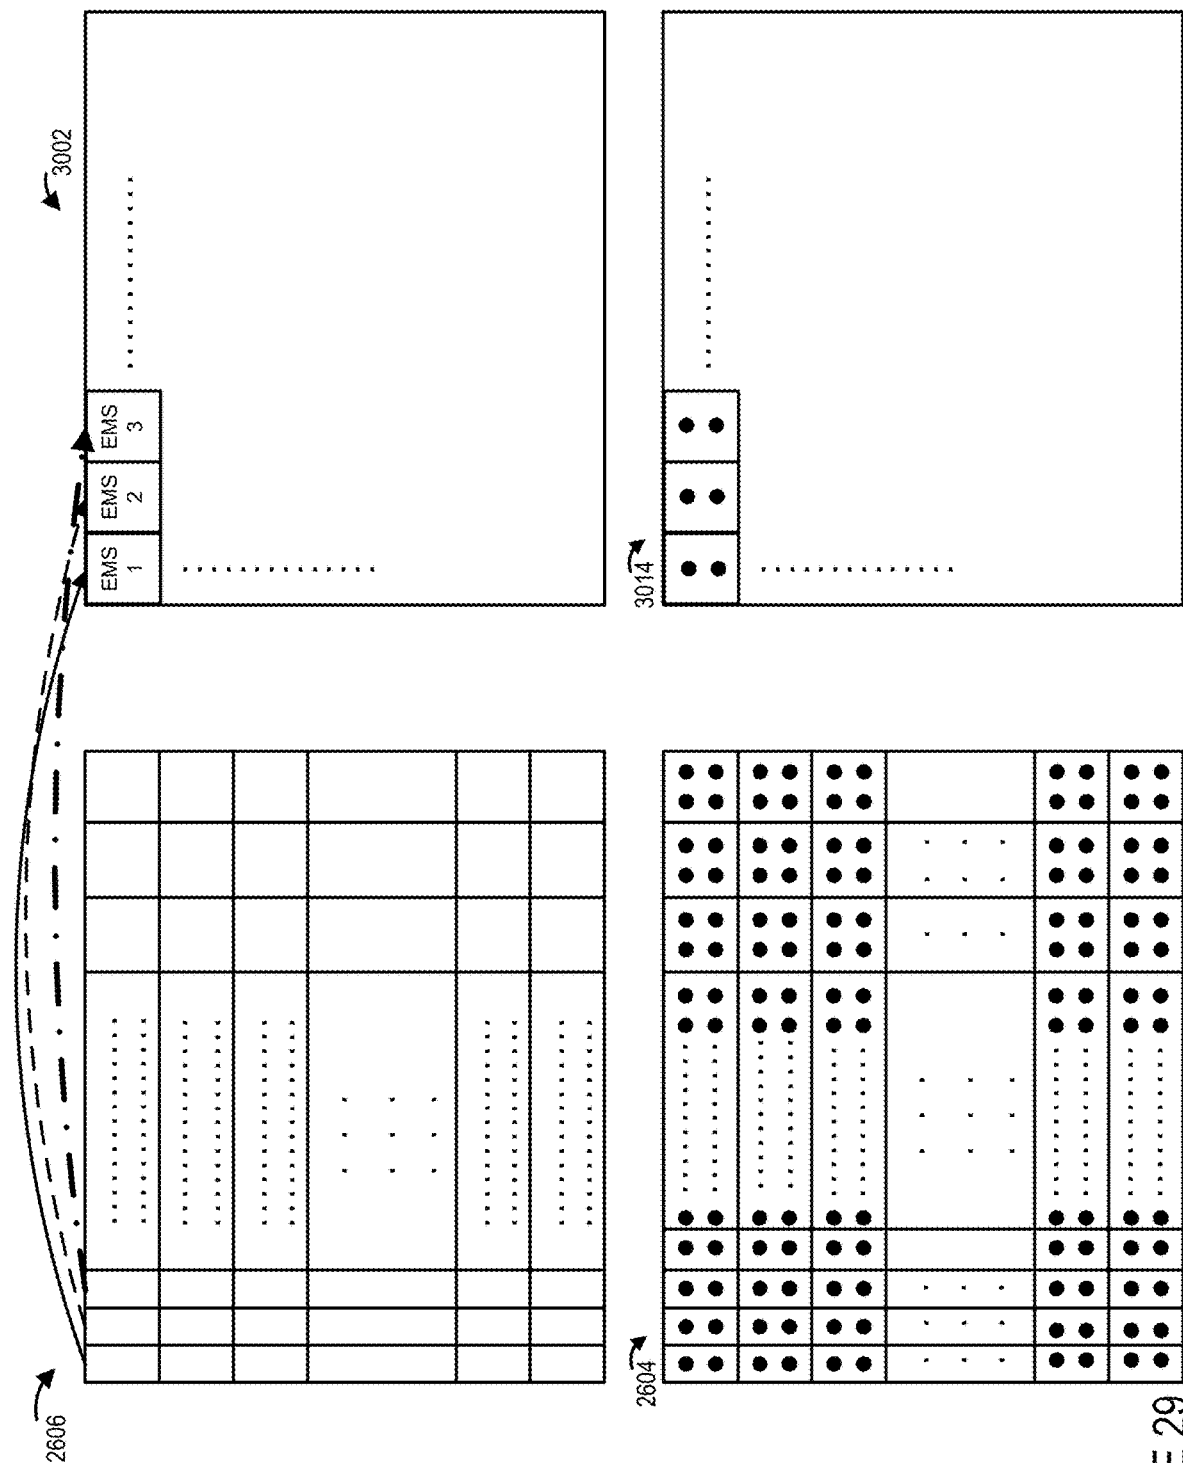


FIGURE 29

METHODS AND APPARATUS FOR RECEIVING AND/OR USING REDUCED RESOLUTION IMAGES

RELATED APPLICATIONS

The present application is a continuation of U.S. patent application Ser. No. 16/035,459 filed Jul. 13, 2018 which is a continuation of U.S. patent application Ser. No. 15/672,236 filed Aug. 8, 2017 which is a continuation of U.S. patent application Ser. No. 15/046,327 filed Feb. 17, 2016 which claims the benefit of U.S. Provisional Application Ser. No. 62/117,427 filed Feb. 17, 2015, U.S. Provisional Application Ser. No. 62/262,374 filed Dec. 2, 2015, and U.S. Provisional Application Ser. No. 62/296,065 filed Feb. 16, 2016 with each of the listed applications being hereby expressly incorporated by reference in its entirety.

FIELD

The present invention relates to methods and apparatus for capturing, streaming and/or playback of content, e.g., content which can be used to simulate an environment.

BACKGROUND

Display devices which are intended to provide an immersive experience normally allow a user to turn his head and experience a corresponding change in the scene which is displayed. Head mounted displays sometimes support 360 degree viewing in that a user can turn around while wearing a head mounted display with the scene being displayed changing as the user's head position changes.

With such devices a user should be presented with a scene that was captured in front of a camera position when looking forward and a scene that was captured behind the camera position when the user turns completely around. While a user may turn his head to the rear, at any given time a user's field of view is normally limited to 120 degrees or less due to the nature of a human's ability to perceive a limited field of view at any given time.

In order to support 360 degrees of view, a 360 degree scene may be captured using multiple cameras with the images being combined to generate the 360 degree scene which is to be made available for viewing.

It should be appreciated that a 360 degree view includes a lot more image data than a simple forward view which is normally captured, encoded for normal television and many other video applications where a user does not have the opportunity to change the viewing angle used to determine the image to be displayed at a particular point in time.

Given transmission constraints, e.g., network data constraints, associated with content being streamed, it may not be possible to stream the full 360 degree view in full high definition video to all customers seeking to receive and interact with the content. This is particularly the case where the content is stereoscopic content including image content intended to correspond to both left eye views and right eye views to allow for a 3D viewing effect.

In view of the above discussion it should be appreciated that there is a need for methods and apparatus for supporting encoding and/or streaming of content in a manner which allows an individual user to be supplied with a wide viewing area, so that the playback device has image data available should a user turn his/her head to view a different portion of the environment while satisfying data transmission constraints.

SUMMARY

Methods and apparatus for supporting delivery, e.g., streaming, of video or other content corresponding to an environment are described. In some embodiments the images corresponding to the environment which are communicated to a playback device exceed the area a user can view at a given time so that content is available in the event the user changes his/her viewing angle by, for example, moving his/her head. By providing images for an environmental area larger than that which can be viewed by a user at a given time the playback device has enough information to provide images should the user's viewing angle change without the playback device having to wait for new images or other content corresponding to a portion of the environment which the user was not previously viewing.

In at least some embodiments the environment is represented using a mesh model. Images are captured and encoded into frames. At the playback device the encoded images are decoded and applied to a surface of the environmental model, e.g., as a texture. The mapping of an image to the surface of the environmental model is in accordance with a texture map also sometimes referred to as a UV map. Generally, but not necessarily in all embodiments, a segment of a UV map corresponds to a segment of the 3D mesh model. In the playback device a UV map is applied to the image and the segments of the image are then applied to the corresponding segments of the 3D model as a texture. In this way a UV map can be used to map a portion of a received image onto a corresponding portion of a model of an environment. To achieve a 3D effect this process is used in some embodiments to map images corresponding to a left eye view onto the 3D model with the result being displayed to a user's left eye. An image corresponding to a right eye view is mapped onto the 3D model to generate an image which is displayed to the user's right eye. Differences between the left and right eye views in 3D embodiments result in a user perceiving images in 3D.

In the case of 3D images where data corresponding to left and right eye images is normally communicated, the amount of data required to support 3D viewing can be considerable since data for two images instead of one needs to be communicated to allow for 3D viewing. Unfortunately, bandwidth constraints in many cases may make it difficult to transmit two full resolution images particularly at the higher resolutions viewers are beginning to expect.

A user's ability to detect the quality of an image decreases with regard to portions of an image the user is not directly looking at. In the case of a scene of an environment a user is likely to be focused on viewing the area of action in the environment, e.g., the portion of the environment where a ball is during a sporting game or where the actors are on a stage or within the environment. The methods and apparatus of the present invention take advantage of this fact to selectively allocate resolution to the image being communicated. While reducing the resolution of images which are less likely to be viewed while maintaining the resolution of portions of images corresponding to an environment which are likely to be viewed, it is possible to make efficient use of limited bandwidth available for streaming image data to a playback device.

In various embodiments images of an environment are captured and selective reduction of resolution is applied. The reduction in resolution may be, and sometimes is, applied to portions of an image perceived to correspond to less important portions of an environment. While an environmental model may remain fixed, in various embodiments the reso-

lution reduction applied to captured images may change as the portion of the environment of high importance changes. For example, while at the start of a soccer game the center field may be considered the important area of the environment since that is where the kickoff occurs, as the ball moves to the left end of the field from the viewers perspective the left end may become more important than the other portions of the field. As the ball moves to the right end of the field the right end of the field may be more important from the viewers perspective than the left and center portions where the ball is not located.

In accordance with one embodiment, a resolution allocation is made based on the relative importance of different portions of an environment at a given time with more resolution being allocated to portions of images corresponding to areas of an environment perceived to be of high importance than areas of low importance. The relative importance may be based on motion detected in captured video providing the images being communicated, from user input such as by tracking where users are looking during the capture of images and/or through control of an operator of the encoding and/or streaming system.

In some embodiments a set of different resolution allocations are supported. Down-sampling or another resolution reduction technique is applied to portions of an image the selected resolution allocation indicates are to be subject to resolution reduction while other portions of the image may be left at full resolution or subject to a lesser amount of resolution reduction. A different texture map is used for different resolution allocations. Thus, while the overall size and/or number of bits of an image communicated to a playback device may be, and sometimes is, the same for different resolution allocations, the texture map (UV map) may be and often will be different for different resolution allocations. In this way different UV maps in combination with selective resolution reduction can be used to allocate different amounts of resolution to different portions of an image of an environment depending on which portion of the environment is considered important at a given point in time while the same environmental model is used despite the different allocations of resolution.

Depending on the embodiment, a set of UV maps corresponding to a portion of the simulated environment may be communicated to the playback device with the streaming device then indicating which UV map is to be used for a communicated image or image pair. An image of an entire 360 degree world view including sky and ground images is communicated in a single frame to a playback device in some embodiments. In other embodiments, images to be used as textures for different portions of an environment are communicated as separate frames. For example, an image of the ground may be sent separately, and not updated as frequently, as an image to be used for a 360 degree horizontal portion of the environment while another image may be sent for a sky view. The resolution allocation selection and indication of a corresponding UV map may be, and sometimes is, performed for each portion of the environment included in the stream as a separate image.

The UV maps, also referred to herein as texture maps and sometimes as image maps, are normally communicated to the playback device before they are required for rendering. They can be communicated in the content stream with which they are to be used or sent separately. Once communicated to a playback device, the UV maps can be, and sometimes are, stored. After a UV map is stored the streaming device

can identify the map by communicating a map identifier in the content stream along with the image or images to which the UV map is applied.

Since the resolution allocation is made prior to encoding, the encoding device and/or streaming device normally associates in the image stream the UV map and/or a map identifier with the communicated image or images. In this way the playback device knows which UV map to use when mapping a received image as part of a rendering operation.

Rendered images are displayed to a user with left eye images of an image pair being displayed to a user's left eye and right eye images displayed to a user's right eye.

By using UV maps in combination with selective resolution reduction, different resolutions can be allocated to portions of an environment in a relatively easy to implement manner without requiring changes to be made to the encoder being used to encode images which have been subject to different resolution allocations and without requiring special decoding of images.

While the methods are well suited for 3D applications where left and right eye images are communicated to provide an stereoscopic image pair, the methods may be and sometimes are used for non-stereoscopic embodiments with selective resolution allocation and corresponding UV maps being used in cases where a single image stream is communicated to a playback device, e.g., with the individual images being decoded and rendered onto the environmental map but with the same image being displayed to both eyes of a user. Alternatively, the methods can be used for embodiments where a single stream of images is communicated to the playback device and the playback device uses computational processing to generate a pair of eye views from a single stream of received images, e.g., by receiving an image and generating a left eye image and a different right eye image from the single received image.

Numerous additional methods and embodiments are described in the detailed description which follows.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 illustrates an exemplary system implemented in accordance with some embodiments of the invention which can be used to capture content, stream content, and output content to one or more users playback devices in accordance with any of the embodiments described herein.

FIG. 2A illustrates an exemplary stereoscopic scene, e.g., a full 360 degree stereoscopic scene which has not been partitioned.

FIG. 2B illustrates an exemplary stereoscopic scene which has been partitioned into 3 exemplary scenes in accordance with one exemplary embodiment.

FIG. 2C illustrates an exemplary stereoscopic scene which has been partitioned into 4 scenes in accordance with one exemplary embodiment.

FIG. 3 illustrates an exemplary process of encoding an exemplary 360 degree stereoscopic scene in accordance with one exemplary embodiment.

FIG. 4 illustrates an example showing how an input image portion is encoded using a variety of encoders to generate different encoded versions of the same input image portion.

FIG. 5 illustrates stored encoded portions of an input stereoscopic scene that has been partitioned into 3 portions.

FIG. 6 illustrates the combination of FIGS. 6A and 6B.

FIG. 6A illustrates a first part of a flowchart illustrating the steps of an exemplary method of streaming content in accordance with an exemplary embodiment implemented

5

using the system of FIG. 1 in which selective resolution allocation and different UV maps are used at different times.

FIG. 6B illustrates a second part of a flowchart illustrating the steps of an exemplary method of streaming content in accordance with an exemplary embodiment implemented using the system of FIG. 1 in which selective resolution allocation and different UV maps are used at different times.

FIG. 7 illustrates an exemplary content delivery system with resolution allocation selection, resolution reduction and encoding capability that can be used to encode and stream content, along with corresponding UV maps, in accordance with the features of the invention.

FIG. 8 illustrates an exemplary content playback device that can be used to receive, decode and display the content streamed by the system of FIG. 7 and may use the UV maps shown and described with reference to FIG. 24 and/or various other figures to allow different UV maps to be used for images having different resolution allocations.

FIG. 9 illustrates the combination of FIGS. 9A and 9B.

FIG. 9A illustrates the first part of an exemplary method of operating a content playback device in accordance with the present invention.

FIG. 9B illustrates the second part of an exemplary method of operating a content playback device in accordance with the present invention.

FIG. 10 illustrates an exemplary method of communicating information to be used to represent an environment in accordance with the present invention.

FIG. 11 illustrates an exemplary image capture and content streaming method in accordance with an exemplary embodiment in which different resolution allocations can be used for images corresponding to the same environmental portion at different times.

FIG. 12 illustrates a method of operating a playback device or system, e.g., a rendering device, which can be used in the system of FIG. 1, to receive and render images using UV maps and an environmental model in accordance with one exemplary embodiment.

FIG. 13 illustrates a camera rig including multiple camera pairs for capturing left and right eye images corresponding to different sectors of a 360 degree field of view along with a camera or cameras directed towards the sky to capture a sky view.

FIG. 14 shows how 5 different environmental mesh maps, corresponding to different camera views, can be combined to create a complete spherical view/environment onto which captured images can be projected, e.g., onto the inner surface, as part of a playback operation.

FIG. 15 shows the full assembly of 5 meshes shown in FIG. 14 to create a spherical simulated environment which can be viewed from a user as if he/she were located at the center of the environment, e.g., sphere.

FIG. 16 shows a left eye view image and a right eye view image captured by left and right eye cameras, with fisheye lenses, corresponding to a sector of the camera rig shown in FIG. 13.

FIG. 17A shows an exemplary mesh model of an environment in accordance with the invention.

FIG. 17B shows a UV map which can be used to map portions of a 2D image onto surfaces of the mesh model shown in FIG. 17A.

FIG. 18 shows how captured left and right eye view images of FIG. 16 may appear after cropping prior to encoding and transmission to one or more playback devices.

6

FIG. 19 shows an environmental mesh model corresponding to one sector of the camera rig with one of the images shown in FIG. 18 applied, e.g., projected, onto the environmental mesh.

FIG. 20 shows application of images captured by cameras corresponding to each of the sectors as well as the sky and ground cameras of the camera rig can be combined and projected onto the modeled environment to simulate a complete 360 environment in the form of a sphere.

FIG. 21 shows how selective resolution can be used with regard to a frame which maps to an environmental grid with different resolutions being used for different portions of the image to be mapped to the environmental model, e.g., with smaller portions of the transmitted image being mapped to corresponding portions of the sky and ground mesh segments than the segments of the middle portion of the environment resulting in lower resolution being allocated to the top and bottom portions of the environment than the middle portion of the environment.

FIG. 22 shows a first captured image of a first portion of an environment, a first resolution adjusted image generated using a first resolution allocation from the first captured image, and a first UV map corresponding to the first resolution allocation.

FIG. 23 shows a second captured image of the first portion of the environment, a second resolution adjusted image generated using a second resolution allocation from the second captured image, and a second UV map corresponding to the second resolution allocation.

FIG. 24 shows a third captured image of the first portion of the environment, a third resolution adjusted image generated using a third resolution allocation from the third captured image, and a third UV map corresponding to the second resolution allocation.

FIG. 25 illustrates the combination of FIGS. 25A and 25B.

FIG. 25A shows a first part of a method of operating a content processing and delivery system in accordance with an exemplary embodiment.

FIG. 25B shows a second part of a method of operating a content processing and delivery system in accordance with an exemplary embodiment.

FIG. 26 illustrates an exemplary embodiment of a method of playing back content in accordance with the invention.

FIG. 27 illustrates an example of how a playback device, such as the playback device or devices shown in any of the other figures, can perform image rendering using a UV map corresponding to the resolution allocation that was used to generate the image to be rendered.

FIG. 28 illustrates an example of how a playback device, such as the playback device or devices shown in any of the other figures, can perform image rendering using a UV map corresponding to the resolution allocation that was used to generate the image to be rendered.

FIG. 29 illustrates an example of how a playback device, such as the playback device or devices shown in any of the other figures, can perform image rendering using a UV map corresponding to the resolution allocation that was used to generate the image to be rendered.

DETAILED DESCRIPTION

FIG. 1 illustrates an exemplary system 100 implemented in accordance with some embodiments of the invention. The system 100 supports content delivery, e.g., imaging content delivery, to one or more customer devices, e.g., playback devices/content players, located at customer premises. The

system **100** includes the exemplary image capturing device **102**, a content delivery system **104**, a communications network **105**, and a plurality of customer premises **106**, . . . , **110**. The image capturing device **102** supports capturing of stereoscopic imagery. The image capturing device **102** captures and processes imaging content in accordance with the features of the invention. The communications network **105** may be, e.g., a hybrid fiber-coaxial (HFC) network, satellite network, and/or internet.

The content delivery system **104** includes an image processing, calibration and encoding apparatus **112** and a content delivery device, e.g. a streaming server **114**. The image processing, calibration and encoding apparatus **112** is responsible for performing a variety of functions including camera calibration based on one or more target images and/or grid patterns captured during a camera calibration process. Content delivery device **114** may be implemented as a server with, as will be discussed below, the delivery device responding to requests for content with image calibration information, optional environment information, and one or more images captured by the camera rig **102** which can be used in simulating a 3D environment. Streaming of images and/or content may be and sometimes is a function of feedback information such as viewer head position and/or user selection of a position at the event corresponding to a camera **102** which is to be the source of the images. For example, a user may select or switch between images from a camera rig positioned at center line to a camera rig positioned at the field goal with the simulated 3D environment and streamed images being changed to those corresponding to the user selected camera rig. Thus it should be appreciated that while a single camera rig **102** is shown in FIG. **1** multiple camera rigs may be present in the system and located at different physical locations at a sporting or other event with the user being able to switch between the different positions and with the user selections being communicated from the playback device **122** to the content server **114**. While separate devices **112**, **114** are shown in the image processing and content delivery system **104**, it should be appreciated that the system may be implemented as a single device including separate hardware for performing the various functions or with different functions being controlled by different software or hardware modules but being implemented in or on a single processor.

Encoding apparatus **112** may, and in some embodiments does, include one or a plurality of encoders for encoding image data in accordance with the invention. The encoders may be used in parallel to encode different portions of a scene and/or to encode a given portion of a scene to generate encoded versions which have different data rates. Using multiple encoders in parallel can be particularly useful when real time or near real time streaming is to be supported.

The content streaming device **114** is configured to stream, e.g., transmit, encoded content for delivering the encoded image content to one or more customer devices, e.g., over the communications network **105**. Via the network **105**, the content delivery system **104** can send and/or exchange information with the devices located at the customer premises **106**, **110** as represented in the figure by the link **120** traversing the communications network **105**.

While the encoding apparatus **112** and content delivery server are shown as separate physical devices in the FIG. **1** example, in some embodiments they are implemented as a single device which encodes and streams content. The encoding process may be a 3D, e.g., stereoscopic, image encoding process where information corresponding to left and right eye views of a scene portion are encoded and

included in the encoded image data so that 3D image viewing can be supported. The particular encoding method used is not critical to the present application and a wide range of encoders may be used as or to implement the encoding apparatus **112**.

Each customer premise **106**, **110** may include a plurality of devices/players, e.g., decoding apparatus to decode and playback/display the image content streamed by the content streaming device **114**. Customer premise **1** **106** includes a decoding apparatus/playback device **122** coupled to a display device **124** while customer premise **N** **110** includes a decoding apparatus/playback device **126** coupled to a display device **128**. In some embodiments the display devices **124**, **128** are head mounted stereoscopic display devices.

In various embodiments decoding apparatus **122**, **126** present the imaging content on the corresponding display devices **124**, **128**. The decoding apparatus/players **122**, **126** may be devices which are capable of decoding the imaging content received from the content delivery system **104**, generate imaging content using the decoded content and rendering the imaging content, e.g., 3D image content, on the display devices **124**, **128**. Any of the decoding apparatus/playback devices **122**, **126** may be used as the decoding apparatus/playback device **800** shown in FIG. **8**. A system/playback device such as the one illustrated in FIG. **8** can be used as any of the decoding apparatus/playback devices **122**, **126**.

FIG. **2A** illustrates an exemplary stereoscopic scene **200**, e.g., a full 360 degree stereoscopic scene which has not been partitioned. The stereoscopic scene may be and normally is the result of combining image data captured from multiple cameras, e.g., video cameras, often mounted on a single video capture platform or camera mount.

FIG. **2B** illustrates a partitioned version **250** of the exemplary stereoscopic scene **200** where the scene has been partitioned into 3 (N=3) exemplary portions, e.g., a front 180 degree portion, a left rear 90 degree portion and a right rear 90 degree portion in accordance with one exemplary embodiment.

FIG. **2C** illustrates another portioned version **280** of the exemplary stereoscopic scene **200** which has been partitioned into 4 (N=4) portions in accordance with one exemplary embodiment.

While FIGS. **2B** and **2C** show two exemplary partitions, it should be appreciated that other partitions are possible. For example the scene **200** may be partitioned into twelve (n=12) 30 degree portions. In one such embodiment, rather than individually encoding each partition, multiple partitions are grouped together and encoded as a group. Different groups of partitions may be encoded and streamed to the user with the size of each group being the same in terms of total degrees of scene but corresponding to a different portion of an image which may be streamed depending on the user's head position, e.g., viewing angle as measured on the scale of 0 to 360 degrees.

FIG. **3** illustrates an exemplary process of encoding an exemplary 360 degree stereoscopic scene in accordance with one exemplary embodiment. The input to the method **300** shown in FIG. **3** includes 360 degree stereoscopic image data **302** captured by, e.g., a plurality of cameras arranged to capture a 360 degree view of a scene. The stereoscopic image data **302**, e.g., stereoscopic video, may be in any of a variety of known formats and includes, in most embodiments, left and right eye image data used to allow for a 3D experience. While the methods are particularly well suited for stereoscopic video, the techniques and methods

described herein can also be applied to 2D images, e.g., of a 360 degree or small scene area.

In step **304** the scene data **302** is partitioned into data corresponding to different scene areas, e.g., N scene areas corresponding to different viewing directions. For example, in one embodiment such as the one shown in FIG. **2B** the 360 degree scene area is partitioned into three partitions a left rear portion corresponding to a 90 degree portion, a front 180 degree portion and a right rear 90 degree portion. The different portions may have been captured by different cameras but this is not necessary and in fact the 360 degree scene may be constructed from data captured from multiple cameras before being dividing into the N scene areas as shown in FIG. **2B** and **2C**.

In step **306** the data corresponding to the different scene portions is encoded in accordance with the invention. In some embodiments each scene portion is independently encoded by multiple encoders to support multiple possible bit rate streams for each portion. In step **308** the encoded scene portions are stored, e.g., in the content delivery system **104**, for streaming to the customer playback devices.

FIG. **4** is a drawing **400** illustrating an example showing how an input image portion, e.g., a 180 degree front portion of a scene, is encoded using a variety of encoders to generate different encoded versions of the same input image portion.

As shown in drawing **400**, an input scene portion **402** e.g., a 180 degree front portion of a scene, is supplied to a plurality of encoders for encoding. In the example there are K different encoders which encode input data with different resolutions and using different encoding techniques to generate encoded data to support different data rate streams of image content. The plurality of K encoders include a high definition (HD) encoder **1 404**, a standard definition (SD) encoder **2 406**, a reduced frame rate SD encoder **3 408**, . . . , and a high compression reduced frame rate SD encoder **K 410**.

The HD encoder **1 404** is configured to perform full high definition (HD) encoding to produce high bit rate HD encoded image **412**. The SD encoder **2 406** is configured to perform low resolution standard definition encoding to produce a SD encoded version **2 414** of the input image. The reduced frame rate SD encoder **3 408** is configured to perform reduced frame rate low resolution SD encoding to produce a reduced rate SD encoded version **3 416** of the input image. The reduced frame rate may be, e.g., half of the frame rate used by the SD encoder **2 406** for encoding. The high compression reduced frame rate SD encoder **K 410** is configured to perform reduced frame rate low resolution SD encoding with high compression to produce a highly compressed reduced rate SD encoded version **K 420** of the input image.

Thus it should be appreciated that control of spatial and/or temporal resolution can be used to produce data streams of different data rates and control of other encoder settings such as the level of data compression may also be used alone or in addition to control of spatial and/or temporal resolution to produce data streams corresponding to a scene portion with one or more desired data rates.

FIG. **5** illustrates stored encoded portions **500** of an input stereoscopic scene that has been partitioned into 3 exemplary portions. The stored encoded portions may be stored in the content delivery system **104**, e.g., as data/information in the memory. The stored encoded portions **500** of the stereoscopic scene includes 3 different sets of encoded portions, with each portion corresponding to a different scene area and each set including a plurality of different encoded versions of the corresponding scene portion. Each encoded version is

a version of encoded video data and thus represents multiple frames which have been coded. It should be appreciated that each encoded version **510**, **512**, **516** is video that corresponds to multiple periods of time and that when streaming, the portion, e.g., frames, corresponding to the period of time being played back will be used for transmission purposes.

As illustrated and discussed above with regard to FIG. **4**, each scene portion, e.g., front, rear scene portions, may be encoded using a plurality of different encoders to produce K different versions of the same scene portion. The outputs of each encoder corresponding to a given input scene are grouped together as a set and stored. The first set of encoded scene portions **502** corresponds to the front 180 degree scene portion, and includes encoded version **1 510** of the front 180 degree scene, encoded version **2 512**, . . . , and encoded version **K 516**. The second set of encoded scene portions **504** corresponds to the scene portion **2**, e.g., 90 degree left rear scene portion, and includes encoded version **1 520** of the 90 degree left rear scene portion, encoded version **2 522**, . . . , and encoded version **K 526** of the 90 degree left rear scene portion. Similarly the third set of encoded scene portions **506** corresponds to the scene portion **3**, e.g., 90 degree right rear scene portion, and includes encoded version **1 530** of the 90 degree right rear scene portion, encoded version **2 532**, . . . , and encoded version **K 536** of the 90 degree right rear scene portion.

The various different stored encoded portions of the 360 degree scene can be used to generate various different bit rate streams for sending to the customer playback devices.

The content delivery system **104** can support a large number of concurrent users since, the encoding process allows the N portions of a scene to be transmitted and processed differently to different users without having to encode the content separately for each individual user. Thus, while a number of parallel encoders may be used to support real time encoding to allow for real or near real time streaming of sports or other events, the number of encoders used tends to be far less than the number of playback devices to which the content is streamed.

While the portions of content are described as portions corresponding to a 360 degree view it should be appreciated that the scenes may, and in some embodiments do, represent a flattened version of a space which also has a vertical dimension. The playback device is able to map the scene portions using a model of the 3D environment, e.g., space, and adjust for vertical viewing positions. Thus, the 360 degrees which are discussed in the present application refer to the head position relative to the horizontal as if a user changed his viewing angle left or right while holding his gaze level.

FIG. **6** which comprises FIGS. **6A** and **6B** is a flowchart **600** illustrating the steps of an exemplary method of providing image content, in accordance with an exemplary embodiment. FIG. **6A** illustrates the first part of the flowchart **600**. FIG. **6B** illustrates the second part of flowchart **600**. The method of flowchart **600** is implemented in some embodiments using the capturing system shown in FIG. **1**.

The method **600** commences in start step **602** shown in FIG. **6A**. Operation proceeds from step **602** to step **604**. In step **604**, a captured image is received. Operation proceeds from step **604** to step **606**.

In step **606**, the resolution allocation to be used is selected. The selection may be made for example based on motion. Operation proceeds from step **606** to decision step **608**. In decision step **608**, if a determination is made that the selected resolution is different from the previous resolution

11

allocation then operation proceeds to step 610. Otherwise operation proceeds to step 612.

In step 610 new downsampling and/or filtering information corresponding to the selected resolution allocation used to control resolution reduction is loaded. Operation proceeds from step 610 to step 612.

In step 612, a resolution reduction operation is performed on the received captured image based on the determined resolution allocation to be used. The resolution reduction operation outputs a reduced resolution image 614 with at least some different image portions having different resolutions. Operation proceeds to step 616.

In step 616, the reduced resolution image is encoded using an encoder which supports compression, e.g., entropy encoding, run length encoding, motion vectors and/or other encoding techniques. Operation proceeds from step 616 to step 618.

In step 618, a UV map corresponding to the resolution allocation to be used for rendering the image subjected to determined resolution allocation, e.g., down sampling, is indicated. By specifying the UV map corresponding to the applied resolution allocation and/or by providing a UV map corresponding to the applied resolution allocation the playback device is provided with information which allows the communicated image to be applied to the 3D model of the environment taking into consideration which portions of the transmitted image were downsampled prior to being communicated to the playback device. Operation proceeds from step 618 to decision step 622 shown on FIG. 6B via connection node A 620.

In decision step 622 a determination is made as to whether the UV map corresponding to the applied resolution allocation has been communicated to the playback device. If the determination is that the UV map corresponding to the applied resolution allocation has not been communicated to the playback device then operation proceeds to step 624. If the determination is that the UV map corresponding to the applied resolution allocation has been communicated to the playback device then operation proceeds to step 626.

In step 624, the UV map corresponding to the applied resolution allocation is communicated to the playback device. Operation proceeds from step 624 to step 626.

In step 626, information indicating the UV map to use is communicated to the playback device. Operation proceeds from step 626 to step 628. In step 628, the encoded image is communicated to the playback device. This method may be executed with respect to each received captured image.

FIG. 7 illustrates an exemplary content delivery system 700 with encoding capability that can be used to encode and stream content in accordance with the features of the invention.

The system may be used to perform encoding, storage, and transmission and/or content output in accordance with the features of the invention. In some embodiments the system 700 or the elements therein perform the operation corresponding to the process illustrated in FIG. 6. The content delivery system 700 may be used as the system 104 of FIG. 1. While the system shown in FIG. 7 is used for encoding, processing and streaming of content, it should be appreciated that the system 700 may also include the ability to decode and display processed and/or encoded image data, e.g., to an operator.

The system 700 includes a display 702, input device 704, input/output (I/O) interface 706, a processor 708, network interface 710 and a memory 712. The various components of

12

the system 700 are coupled together via bus 709 which allows for data to be communicated between the components of the system 700.

The memory 712 includes various modules, e.g., routines, which when executed by the processor 708 control the system 700 to implement the partitioning, encoding, storage, and streaming/transmission and/or output operations in accordance with the invention.

The memory 712 includes various modules, e.g., routines, which when executed by the processor 707 control the computer system 700 to implement the immersive stereoscopic video acquisition, encoding, storage, and transmission and/or output methods in accordance with the invention. The memory 712 includes control routines 714, a partitioning module 716, encoder(s) 718, a streaming controller 720, received input images 732, e.g., 360 degree stereoscopic video of a scene, encoded scene portions 734, and timing information 736. In some embodiments the modules are implemented as software modules. In other embodiments the modules are implemented in hardware, e.g., as individual circuits with each module being implemented as a circuit for performing the function to which the module corresponds. In still other embodiments the modules are implemented using a combination of software and hardware.

The control routines 714 include device control routines and communications routines to control the operation of the system 700. The partitioning module 716 is configured to partition a received stereoscopic 360 degree version of a scene into N scene portions in accordance with the features of the invention.

The encoder(s) 718 may, and in some embodiments do, include a plurality of encoders configured to encode received image content, e.g., 360 degree version of a scene and/or one or more scene portions in accordance with the features of the invention. In some embodiments encoder(s) include multiple encoders with each encoder being configured to encode a stereoscopic scene and/or partitioned scene portions to support a given bit rate stream. Thus in some embodiments each scene portion can be encoded using multiple encoders to support multiple different bit rate streams for each scene. An output of the encoder(s) 718 is the encoded scene portions 734 which are stored in the memory for streaming to customer devices, e.g., playback devices. The encoded content can be streamed to one or multiple different devices via the network interface 710.

UV maps 740 are stored in memory 712 of the content delivery system 700. The UV maps 740 correspond to different resolution allocations and/or areas of the environment. For example, the first UV map 1 742 corresponds to a first resolution allocation, the second UV map 2 744 corresponds to a second resolution allocation, and the third UV map 746 corresponds to a third resolution allocation. UV maps with different resolution allocations can correspond to the same area of an environment. Different UV maps corresponding to other areas of the environment can be stored in the memory 712. Multiple UV maps may correspond to the environmental model. The mesh model of the environment where the received images were captured is stored in memory 712 of the content delivery system 700, e.g., 3D environmental mesh model 738. Multiple mesh models may be stored in the memory 712.

The streaming controller 720 is configured to control streaming of encoded content for delivering the encoded image content to one or more customer devices, e.g., over the communications network 105. In various embodiments various steps of the flowchart 600 are implemented by the

13

elements of the streaming controller 720. The streaming controller 720 includes a request processing module 722, a data rate determination module 724, a current head position determination module 726, a selection module 728 and a streaming control module 730. The request processing module 722 is configured to process a received request for imaging content from a customer playback device. The request for content is received in various embodiments via a receiver 713 in the network interface 710. In some embodiments the request for content includes information indicating the identity of requesting playback device. In some embodiments the request for content may include data rates supported by the customer playback device, a current head position of the user, e.g., position of the head mounted display. The request processing module 722 processes the received request and provides retrieved information to other elements of the streaming controller 720 to take further actions. While the request for content may include data rate information and current head position information, in various embodiments the data rate supported by the playback device can be determined from network tests and other network information exchange between the system 700 and the playback device.

The data rate determination module 724 is configured to determine the available data rates that can be used to stream imaging content to customer devices, e.g., since multiple encoded scene portions are supported the content delivery system 700 can support streaming content at multiple data rates to the customer device. The data rate determination module 724 is further configured to determine the data rate supported by a playback device requesting content from system 700. In some embodiments the data rate determination module 724 is configured to determine data rates for delivery of image content based on network measurements.

The current head position determination module 726 is configured to determine a current viewing angle and/or a current head position of the user, e.g., position of the head mounted display, from information received from the playback device. In some embodiments the playback device periodically sends current head position information to the system 700 where the current head position determination module 726 receives and processes the information to determine the current viewing angle and/or a current head position.

The selection module 728 is configured to determine which portions of a 360 degree scene to stream to a playback device based on the current viewing angle/head position information of the user. The selection module 728 is further configured to select the encoded versions of the determined scene portions based on the available data rates to support streaming of content.

The streaming control module 730 is configured to control streaming of image content, e.g., multiple portions of a 360 degree stereoscopic scene, at various supported data rates in accordance with the features of the invention. In some embodiments the streaming control module 730 is configured to control the streaming of N portions of a 360 degree stereoscopic scene to the playback device requesting content to initialize scene memory in the playback device. In various embodiments the streaming control module 730 is configured to send the selected encoded versions of the determined scene portions periodically, e.g., at a determined rate. In some embodiments the streaming control module 730 is further configured to send 360 degree scene updates to the playback device in accordance with a time interval, e.g., once every minute. In some embodiments sending 360 degree scene update includes sending N scene portions or

14

N-X scene portions of the full 360 degree stereoscopic scene, where N is the total number of portions into which the full 360 degree stereoscopic scene has been partitioned and X represents the selected scene portions recently sent to the playback device. In some embodiments the streaming control module 730 waits for a predetermined time after initially sending N scene portions for initialization before sending the 360 degree scene update. In some embodiments the timing information to control sending of the 360 degree scene update is included in the timing information 736. In some embodiments the streaming control module 730 is further configured identify scene portions which have not been transmitted to the playback device during a refresh interval; and transmit an updated version of the identified scene portions which were not transmitted to the playback device during the refresh interval.

In various embodiments the streaming control module 730 is configured to communicate at least a sufficient number of the N portions to the playback device on a periodic basis to allow the playback device to fully refresh a 360 degree version of said scene at least once during each refresh period.

FIG. 8 illustrates a computer system/playback device 800 implemented in accordance with the present invention which can be used to receive, decode, store and display imaging content received from a content delivery system such as the one shown in FIGS. 1 and 7. The playback device may be used with a 3D head mounted display such as the OCULUS RIFT™ VR (virtual reality) headset which may be the head mounted display 805. The device 800 includes the ability to decode the received encoded image data and generate 3D image content for display to the customer. The playback device in some embodiments is located at a customer premise location such as a home or office but may be located at an image capture site as well. The device 800 can perform signal reception, decoding, display and/or other operations in accordance with the invention.

The device 800 includes a display 802, a display device interface 803, input device 804, a decoder 864, input/output (I/O) interface 806, a processor 808, network interface 810 and a memory 812. The various components of the playback device 800 are coupled together via bus 809 which allows for data to be communicated between the components of the system 800. While in some embodiments display 802 is included as an optional element as illustrated using the dashed box, in some embodiments an external display device 805, e.g., a head mounted stereoscopic display device, can be coupled to the playback device via the display device interface 803. In some embodiments, the network interface 810 includes a receiver 860 and a transmitter 862.

The memory 812 includes various modules, e.g., routines, which when executed by the processor 808 control the playback device 800 to perform decoding and output operations in accordance with the invention. The memory 812 includes control routines 814, a request for content generation module 816, a head position and/or viewing angle determination module 818, a decoder module 820, a stereoscopic image rendering module 822 also referred to as a 3D image generation module, and data/information including received encoded image content 824, decoded image content 826, a 360 degree decoded scene buffer 828, and generated stereoscopic content 830.

The control routines 814 include device control routines and communications routines to control the operation of the device 800. The request generation module 816 is configured to generate a request for content to send to a content delivery system for providing content. The request for

15

content is sent in various embodiments via the network interface **810**. The head position and/or viewing angle determination module **818** is configured to determine a current viewing angle and/or a current head position of the user, e.g., position of the head mounted display, and report the determined position and/or viewing angle information to the content delivery system **700**. In some embodiments the playback device **800** periodically sends current head position information to the system **700**.

The decoder module **820** is configured to decode encoded image content **824** received from the content delivery system **700** to produce decoded image data **826**. The decoded image data **826** may include decoded stereoscopic scene and/or decoded scene portions.

The 3D image rendering module **822** generates 3D images in accordance with the features of the invention, e.g., using the decoded image content **826**, for display to the user on the display **802** and/or the display device **805**. The generated stereoscopic image content **830** is the output of the 3D image generation module **822**. Thus the rendering module **822** renders the 3D image content **830** to the display. In some embodiments the display device **805** may be a 3D display such as an oculus rift. The operator of the playback device **800** may control one or more parameters via input device **804** and/or select operations to be performed, e.g., select to display 3D scene.

FIG. 8 illustrates an exemplary content playback device that can be used to receive, decode and display the content streamed by the system of FIG. 7. The system **800** includes a display interface **803** coupled to a head mounted stereoscopic display **805**, an input device **804**, an optional display **802** and I/O interface. The interface **802** coupled the various input/output elements **803**, **802**, **804** to the bus **809** which in turn is coupled to processor **808**, network interface **810** and memory **812**. The network interface **810** allows the playback device to receive content from the streaming device **114** and/or communicate information such as view head position and/or position (camera rig) selection indicating selection of particular viewing position at an event. The memory **812** includes various data and modules as shown in FIG. 8. When executed the decoder module **820** causes received images to be decoded while 3D image rendering module **822** causes further processing of the images in accordance with the present invention and optionally stitching of images together as part of the presentation process.

FIG. 9 which comprises a first part FIG. 9A and a second part FIG. 9B illustrates the steps **900** of a method of operating a content playback device. In accordance with the method **900** different UV maps may be used at different times for mapping a portion of one or more received images to an environmental model, e.g., a mesh model, of an environment. As a result of using different UV maps, while the number of pixels in a received image, e.g., encoded frame, may remain the same, the mapping of pixels of a received image to a segment of the environmental model may change. For example, using a first UV map may result in a first number of pixels in a received image mapping to a first portion of an environmental model while use of a second different UV map may result in a different number of pixels in a received image mapping to the same portion of the environmental model. The system generating and transmitting the images also in some embodiments communicates the UV maps and/or indicates to the playback device which UV map is to be used when mapping an image or set of images to the environmental model. Thus by changing the UV map to be used the encoding and transmission device can change the amount of data and/or resolution associated

16

with a particular portion of the environmental model. Since the rendering involves stretching or otherwise conforming the indicated portion of an image to the corresponding segment of the 3D environmental model the image content will be scaled and/or otherwise modified as needed as part of the rendering process to cover the segment of the 3D model to which it applies. Consider for example if a first UV map maps one pixel to a first segment of the environmental model and a second UV map maps two pixels to the first segment of the environmental model, the resolution of the displayed first segment will be higher when the second UV map is used than when the first UV map is used for image rendering. While the UV map may be changed from image to image or from group of images to group of images thereby allowing the server generating and sending the images and UV map information to the playback device to dynamically alter the allocation of data and/or resolution within a portion of the environment, e.g., front portion, based on the scene areas considered of particular interest, e.g., scene areas where the actors, players, performers are in the environment or where movement is in the environment, the data rate used for transmitting images can be held relatively constant since the number of pixels in the images can remain the same with the UV map controlling the allocation of pixels to portions of the environment. Thus the methods allow for the image encoding technique to remain the same at least in some embodiments with the captured image or images being downsampled differently prior to encoding depending on the location of the scene portions considered of particular interest within a captured image and based on knowledge of which UV map will be used to apply the image, e.g., as a texture, to one or more segments of an environmental module. While the UV map may be changed on a per frame or image basis from one image or frame to the next, in some embodiments the change in UV maps is constrained to occur on I-frame or group of picture boundaries with a UV map being used for multiple frames within a group of pictures or between I-frames. While such a UV map transition constraint is used in some embodiments, it is not necessary or critical to the invention and some embodiments allow the UV map to be changed on a per frame basis.

The steps of the exemplary method **900** will now be discussed in detail. The method **900** starts in step **902**, e.g., with a content playback device being powered on. The playback device may be, e.g., a game system connected to a head mounted display or TV or as is the case in various embodiments a cell phone mounted in a head mount with a touch pad or other control and one or more lenses for allowing a user to view left and right eye images on different portions of a cell phone screen which is used as a display device. The method **900** may be implemented by any of the content playback devices described in the present application.

In step **903**, e.g., in response to user input indicating user selection of content to be played to a user, the content playback device transmits a request for content in step **903**. In some embodiments this request is communicated to a content server or content provider system, e.g., a device which receives, processes and encodes images of an environment and supplies them to the playback device along with UV maps and/or information about which UV map to be used at a given time. The server may also provide an environmental model or a default model may be used.

In step **904** a model of an environment, e.g., a 3D mesh model is received, e.g., from the content server. The model may be and sometimes is a model of an environment where an event such as a play or sporting event is ongoing. The

model may be a complete 360 degree model of the environment or a model of the portion of the environment to which image content is to be mapped, e.g., a front portion of the environment. As should be appreciated the features relating to using different UV maps to map images to a portion of the environment may be used for a full 360 degree environment, a portion of an environment, with stereo images and/or with non-stereoscopic images, e.g., panoramic images where the same image is displayed to both left and right eyes of a viewer.

Operation proceeds from step 904 to step 906 in which model of the environment received in step 903 is stored for future use, e.g., in rendering and displaying images mapped onto the model in accordance with one of the UV maps, e.g., texture maps, which are received in step 908. The texture maps may be and sometimes are received from the same server which provides the environmental model. The UV map indicates how a 2d image should be segmented with the segments then being applied to corresponding segments of the environmental model, e.g., as a texture or textures.

While an initial texture, e.g., initial UV map, may be received in step 908 in some embodiments a set of maps are received and stored with the different UV maps indicating different mappings between an image and a portion of the environmental model. Each map may be identified by a texture map identifier. During streaming of content the content server providing the images can indicate which texture map to use with which set of images. In other embodiments a new texture map may be streamed with or before the images to which the new texture map is to be applied. Storing of a set of texture maps in the playback device can provide efficient transmission since the maps can be reused without transmitting the UV/texture maps multiple times to the playback device.

In step 910 the received set of texture maps is stored for future use. With the texture maps having been stored, operation proceeds to step 914 in which image content is received. In step 904 in addition to image content an indicator identifying the texture map to be used to map the received image onto the model of the environment is received or the texture map to be used is received. When an indicator is received it identifies the texture map in the stored set of texture maps which is to be used. An indicated texture map may remain in effect until a new texture map is specified and/or provided. Thus a single texture map may be used for a sequence of images, e.g., a group of pictures. The texture map may be changed by the server when, e.g., motion is detected indicating a different area of the environment is an area of higher priority than an area to which high resolution was previously allocated. Thus as actors move or players on a field move, resolution allocation can be changed and the UV map corresponding to the current resolution allocation may be used in place of a previous UV map corresponding to a different resolution allocation.

Step 914 includes, in some embodiments steps 916, 918, 920, 926 and 928.

In step 916 a first encoded image is received. In step 918 which is optional, a second encoded image is received.

In step 920 which is an alternative to steps 916, 918 an encoded frame including one or both images is received. The second encoded image may be a second image of a stereoscopic image pair with the first and second images being left and right eye images to be displayed to a user of the playback device. For example odd lines of a frame may provide the first image and even lines of the encoded frame may provide the second encoded image. Alternatively a top half of an encoded frame may provide the first image and the

bottom half the second image. Other ways of including the first and second images in a single frame are also possible.

In step 914, in addition to receiving image content which can be mapped to the environmental model, in step 926 a first indicator indicating which of a plurality of texture maps corresponding to different resolution allocation is to be used with the received first and/or second encoded images is also received. If a new texture map indicator is not received in step 914, and a new texture map is not received, the playback device will continue to use the last UV map which was being used. Rather than receive a texture map indicator, a new texture map may be received in step 928 which is to be used in rendering the received images.

With images received, e.g., in encoded form, operation proceeds from step 914 to step 930. In step 930 the received image or images are decoded. For example in step 932 the first encoded image is decoded to recover a first image. In step 934 the second encoded image is decoded to recover a second image. As discussed above, the first and second images may be left and right eye views. In embodiments where the first and second images are included in a single encoded frame decoding of the received frame and separation of the left and second images may be used in step 930 to produce left and right eye images which may be and sometimes are applied separately to the environmental map to generate separate, potentially different, left and right eye views.

In some embodiments the images communicate a complete 360 degree environment or panoramic view. In other embodiments the first and second images may correspond to a portion of the environment, e.g., a front portion or a 360 degree middle panoramic portion but not the sky and ground. In step 936 other images which are sent with the encoded first and second image or in a separate stream may be decoded to obtain textures for portions of the environment which are not provided by the first and/or second images. In some embodiments in step 936 a sky or ground image is obtained by decoding a received encoded image or frame.

With the decoding of images that were transmitted to the playback device having been completed in step 930 operation proceeds to step 938 in which image content is rendered using the received, e.g., decoded image or images, the UV map which was to be used in rendering the received images, and the environmental model. Step 938 involves applying the first image to the environmental model in accordance with UV map to be used. Thus the first image is used as a texture which is applied to segments of the environmental model in accordance with the applicable UV map, e.g., a first UV map. The rendering may be performed separately for left and right eye views.

In some embodiments step 938 includes step 940. In step 940 the first image is rendered by using the first texture map (UV map) corresponding to a first resolution allocation to apply at least a portion of the first image to a surface of a first portion, e.g., first segment, of the model of the environment. For example a first set of pixels of the first image may be mapped to the first segment of the mesh model of the environment based on the first texture map. In step 942 which may be performed in the case of stereo image playback, the second image is rendered by using the first texture map (UV map) corresponding to a first resolution allocation to apply at least a portion of the second image to a surface of the first portion, e.g., the first segment, of the model of the environment. For example a first set of pixels of the second image may be mapped to the first segment of the mesh model of the environment based on the first texture

map. In optional step **944** images of portions of the environment not included in the first image, e.g., the sky or ground portions, are rendered, e.g., applied to the environmental model in accordance with a UV map relevant to these portions. It should be appreciated in some embodiments separate sky and ground portions are not communicated with such portions being part of the first and second images in some embodiments.

In operation step **946**, which is performed for each eye view, rendered images corresponding to different portions of a 360 degree simulated environment are combined to the extent need to provide a contiguous viewing area to a user. Step **946** is performed separately for the left and right eye images since while the ground and sky portions may be the same for these images when they are presented in non-stereo format, the other portions of the left and right eye images may include differences which may result in the perception of depth when the left and right eye images are viewed by different eyes of a user of the playback device.

With the first image or pair of images having been applied to the model of the environment in step **938** operation proceeds to display step **950** via connecting node **948**. In step **950** the rendered image content is displayed to a user of the playback device, e.g., on a display screen. In step **952** a first rendered image or a combined image generated from recovered first image content is displayed for viewing by one of a users left and right eyes or both eyes if stereoscopic display is not supported. In step **954** which is performed in the case of stereoscopic display, a second rendered image is displayed to a second one of a users left and right eyes. The displayed rendered second image is an image that was generated from recovered second image data or a combination of recovered, e.g., decoded second image data and data from another image, e.g., a sky or ground image portion.

With one image or pair of images having been rendered and displayed, operation proceeds to step **956** in which content corresponding to another image or pair of images is received and processed. The image or images received in step **956** may be and sometimes do correspond to a second group of pictures and corresponds to a different point in time than the first image. Thus, between the time the first image was captured and the third image received in step **956** was captured the players, actors or an area of motion may have shifted position from where the activity was at the time the first image was captured. For example, while remaining in a forward field of view, the players on a field may have moved left triggering the sever providing the third image to use a resolution allocation giving more resolution to the left portion of the front field of view than a center or right portion where the action was at the time the first image was captured. The different resolution allocation, e.g., a second resolution allocation by the server or encoding device, will correspond to specification that the playback device should use a different UV map, e.g., a second texture map, for rendering the third image than the first image. For example, the second UV map may specify using fewer pixels from the third image to map to the first segment than were used to map from the first image to the first segment of the environmental map and to use more pixels from third image to map to a second segment located in the left side of the forward field of view in the environmental model where the action is now located at the time of capture of the third image than were used to map to the second segment of the environmental map from the first frame.

Step **956** will now be discussed in detail. Step **956** includes, in some embodiments, steps **958**, **960**, **962**, **926** and/or **968**.

In step **958** a third encoded image is received. In step **960** which is optional, a fourth encoded image is received.

In step **962** which is an alternative to steps **958**, **960** an encoded frame including one or both of the third and fourth encoded images is received.

The third encoded image may be a first image of a second stereoscopic image pair with the third and fourth images being left and right eye images to be displayed to a user of the playback device.

In step **956**, in addition to receiving image content which can be mapped to the environmental model, in step **968** a second indicator indicating which of a plurality of texture maps corresponding to different resolution allocation is to be used with the received third and/or fourth encoded images is also received. If a new texture map indicator is not received, in step **968** and a new texture map is not received, the playback device will continue to use the last UV map which was being used. Rather than receive a texture map indicator, a new texture map may be received in step **970** which is to be used in rendering the received third and/or fourth images.

With images received, e.g., in encoded form, operation proceeds from step **956** to step **970**. In step **970** the received third and/or fourth image or images are decoded. For example in step **974** the third encoded image is decoded to recover a third image. In step **976** the fourth encoded image is decoded to recover a fourth image. As discussed above, the third and fourth images may be left and right eye views. In embodiments where the third and fourth images are included in a single encoded frame decoding of the received frame and separation of the third and fourth images may be performed in step **972** to produce left and right eye images which may be and sometimes are applied separately to the environmental map to generate separate, potentially different, left and right eye views.

In some embodiments the third and/or fourth images communicate a complete 360 degree environment or panoramic view. In other embodiments the third and fourth images may correspond to a portion of the environment, e.g., a front portion or a 360 degree middle panoramic portion but not the sky and ground. In step **978** other images which are sent with the encoded third and fourth images or in a separate stream may be decoded to obtain textures for portions of the environment which are not provided by the third and/or fourth images. In some embodiments in step **986** a sky or ground image is obtained by decoding a received encoded image or frame.

With the decoding of images that were transmitted to the playback device having been completed in step **972** operation proceeds to step **980** in which image content is rendered using the received, e.g., decoded image or images, the UV map which was to be used in rendering the received images, e.g., the second UV map, and the environmental model. Step **980** involves applying the third image to the environmental model in accordance with second UV map to be used which results in a different allocation of pixels from the received image to the model of the environment than occurred when using the first UV map. Thus, as part of the rendering the third image is used as a texture which is applied to segments of the environmental model in accordance with the applicable UV map, e.g., the second UV map. The rendering may be performed separately for left and right eye views.

In some embodiments step **980** includes step **982**. In step **982** the third image is rendered by using the second texture map (UV map) corresponding to a second resolution allocation to apply at least a portion of the third image to a surface of the first portion, e.g., first segment, of the model of the environment. For example a first set of pixels of the

21

third image may be mapped to the first segment of the mesh model of the environment based on the second texture map where the first set of pixels includes fewer pixels than the first set which were mapped when the first UV map was used. A second set of pixels may be mapped to a second segment of the model where the second set of pixels includes more pixels than were mapped to the second segment when the first UV map was used. Thus by using different UV maps to map an image to the model, different allocations of the limited number of pixels to portions of the model of the environment may be achieved in an easy manner without having to alter the number of pixels transmitted in the encoded images provided to the playback device.

In step 978 which may be performed in the case of stereo image playback, the fourth image is rendered by using the second texture map (UV map) corresponding to the second resolution allocation to apply at least a portion of the fourth image to a surface of the first portion, e.g., the first segment, of the model of the environment. Similarly the second UV map is used to control mapping of pixels from the fourth image to the second segment of the environmental model.

In optional step 986 images of portions of the environment not included in the first image, e.g., the sky or ground portions, are rendered, e.g., applied to the environmental model in accordance with a UV map relevant to these portions. It should be appreciated in some embodiments separate sky and ground portions are not communicated with such portions being part of the first and second images in some embodiments.

In operation step 988, which is performed for each eye view, rendered images corresponding to different portions of a 360 degree simulated environment are combined to the extent need to provide a contiguous viewing area to a user. Step 988 is performed separately for the left and right eye images since while the ground and sky portions may be the same for these images when they are presented in non-stereo format, the other portions of the left and right eye images may include differences which may result in the perception of depth when the left and right eye images are viewed by different eyes of a user of the playback device.

With the third image, which may be part of a second pair of images, having been applied to the model of the environment in step 980 operation proceeds to display step 990. In step 990 the rendered image content is displayed to a user of the playback device, e.g., on a display screen. In step 992 a third rendered image or a combined image generated from recovered third image content is displayed for viewing by one of a users left and right eyes or both eyes if stereoscopic display is not supported. In step 994 which is performed in the case of stereoscopic display, a fourth rendered image is displayed to a second one of a user's left and right eyes. The displayed rendered fourth image is an image that was generated from recovered fourth image data or a combination of recovered, e.g., decoded fourth image data and data from another image, e.g., a sky or ground image portion.

The process of receiving and decoding images and rendering images using the UV map provided or specified by the server providing the images occurs on an ongoing basis as represented with operation proceeding from step 990 back to step 914 via connecting node B 996 allowing for additional images to be received and processed, e.g., a new first and second images.

In some embodiments the images correspond to a live sporting event with the server providing the images specifying different UV maps to be used during different portions of the sporting event based on where the action is occurring on the sports field with the generation of the images to be

22

transmitted in encoded form taking into consideration the UV map which will be used to render the images. Thus, by specifying the use of different UV maps at different times resolution can be dynamically allocated to match where the action is occurring on a sports field or in an environment.

FIG. 11 which illustrates an image capture and content streaming method in accordance with an exemplary embodiment. The method 1100 shown in FIG. 11 starts in step 1102 when it is time to capture images, e.g., images corresponding to an event such as a sporting event or music performance.

From start step 1102 operation proceeds along a plurality of paths, the paths bringing with steps 1114, 1104, 1106, 1108, 1110, 1112, which may be performed in parallel and, optionally, asynchronously.

To facilitate an understanding of the image capture process reference will now be made to the exemplary camera rig shown in FIG. 13. The camera rig 1300 can be used as the rig 102 of the FIG. 1 system and includes a plurality of stereoscopic pairs each corresponding to a different one of three sectors. The first camera pair 1301 includes a left eye camera 1302 and a right camera 1304 intended to capture images corresponding to those which would be seen by a left and right eye of a person. Second sector camera pair 1305 includes left and right cameras 1306, 1308 while the third sector camera pair 1309 includes left and right cameras 1310, 1312. Each camera is mounted in a fixed position in the support structure 1318. An upward facing camera 1314 is also included. A downward facing camera which is not visible in FIG. 13 may be included below camera 1314. Stereoscopic camera pairs are used in some embodiments to capture pairs of upward and downward images however in other embodiments a single upward camera and a single downward camera are used. In still other embodiments a downward image is captured prior to rig placement and used as a still ground image for the duration of an event. Such an approach tends to be satisfactory for many applications given that the ground view tends not to change significantly during an event.

The output of the cameras of the rig 1300 are captured and processed by the method of FIG. 11 which will now be discussed further. Image capture steps shown in FIG. 11 are normally performed by operating a camera of the camera rig 102 to capture an image while encoding of images is performed by encoder 112 with responses to streaming requests and streaming of content being preformed by the streaming server 114.

In the first path of FIG. 11 which relates to downward image capture and processing, in step 1114 an image is captured of the ground, e.g., beneath rig 102. This may happen prior to rig placement or during the event if the rig includes a downward facing camera. From step 1114 operation proceeds to steps 1144 where the captured image is cropped prior to encoding in step 1145. The encoded ground image is then stored pending a request for content which may be responded to by supplying one or more encoded images in step 1146 to a requesting device.

The second processing path shown in FIG. 11 which starts with step 1104 relates the processing and responding to requests for content. In step 1104 monitor for request for content occurs, e.g., by content server 114. In step 1128 a request for content is received from a playback device, e.g. device 122 located at customer premise 106.

In response to the content request the playback device is provided with information including one or UV maps corresponding to different resolution allocations which may be used.

23

From step 1104 operation proceeds to step 1128 in which is performed in cases where an environmental map was generated and/or other environmental information which may be different from a predetermined default setting or environmental is supplied to the playback device to be used in rendering images as part of an environmental simulation.

Thus, via step 1132 a playback device requesting content is provided the information need to model the environment and/or with other information which may be needed to render images onto the model. In addition to model information step 1132 may optionally include communication of a set of UV maps to the playback device requesting content for future use, e.g., with some different UV maps corresponding to different resolution allocations but the same area of a model in some embodiments.

In some embodiments when the FIG. 13 camera rig is used each of the sectors corresponds to a known 120 degree viewing area with respect to the camera rig position, with the captured images from different sector pairs being seamed together based on the images known mapping to the simulated 3D environment. While a 120 degree portion of each image captured by a sector camera is normally used, the cameras capture a wider image corresponding to approximately a 180 degree viewing area. Accordingly, captured images may be subject to masking in the playback device as part of the 3D environmental simulation or cropping prior to encoding. FIG. 14 is a composite diagram 1400 showing how a 3D spherical environment can be simulated using environmental mesh portions which correspond to different camera pairs of the rig 102. Note that one mesh portion is shown for each of the sectors of the rig 102 with a sky mesh being used with regard to the top camera view and the ground mesh being used for the ground image captured by the downward facing camera.

When combined the overall meshes corresponding to different cameras results in a spherical mesh 1500 as shown in FIG. 15. Note that the mesh 1500 is shown for a single eye image but that it is used for both the left and right eye images in the case of stereoscopic image pairs being captured.

Mesh information of the type shown in FIG. 14 can and sometimes is communicated to the playback device in step 1132. The communicated information will vary depending on the rig configuration. For example if a larger number of sectors were used masks corresponding to each of the sectors would correspond to a small viewing area than 120 degrees with more than 3 environmental grids being required to cover the diameter of the sphere.

Environmental map information is shown being optionally transmitted in step 1132 to the playback device. It should be appreciated that the environmental map information is optional in that the environment may be assumed to be a default size sphere in the event such information is not communicated having a predetermined number of segments arranged in a known mesh. In cases where multiple different default size spheres are supported an indication as to what size sphere is to be used may be and sometimes is communicated in step 1132 to the playback device.

Operation proceeds from step 1132 to streaming step 1146. Image capture operations may be performed on an ongoing basis during an event particularly with regard to each of the 3 sectors which can be captured by the camera rig 102. Accordingly, processing paths starting with steps 1106, 1108 and 1110 which correspond to first, second and third sectors of the camera rig are similar in terms of their content.

In step 1106, the first sector pair of cameras is operated to capture images, e.g., a left eye image in step 1116 and a right

24

eye image in step 1118. FIG. 16 shows an exemplary image pair 1600 that may be captured in step 1106. The captured images are then cropped in step 1134, e.g., to remove undesired image portions such as image portions captured by another camera pair. In step 1144 a resolution allocation to be used for the captured left and right eye image is determined, e.g., selected. The selection may be based in information about which portion of the environment and thus captured images was important at the time of the capture of the images. The importance information may be based on detection of where individuals at the event being videoed are looking at the time of image capture, system controller input and/or the location of motion in the environment at the time of image capture. A resolution reduction operation is performed on the captured images in step 1146 based on the determined, e.g., selected, resolution allocation. The selected resolution allocation may be one of a plurality of supported resolution allocations corresponding to different supported UV maps corresponding to the portion of the environment captured by the first sector camera pair. In step 1148 the reduced resolution images generated in step 1146 are encoded. Information indicating the UV map to be used for rendering the reduced resolution images generated in step 1146 is generated in step 1149 and will, in some embodiments be associated with and transmitted with the encoded images generated in step 1146 so that the playback device can determine which UV map to use when rendering images recovered by decoding the encoded images generated in step 1146.

FIG. 17A shows an exemplary mesh model 1700 of an environment in accordance with the invention.

FIG. 17B shows a UV map 1702 which can be used to map portions of a 2D image onto surfaces of the mesh model shown in FIG. 17A.

FIG. 18 shows an exemplary result 2000 of cropping the left and right eye view images of FIG. 16 as may occur in step 1134. The cropping of image pair shown in FIG. 18 may be performed prior to encoding and transmission to one or more playback devices.

The image capture, cropping and encoding is repeated on an ongoing basis at the desired frame rate as indicated by the arrow from step 1149 back to step 1106.

Similar operations to those described with regard to the images captured for the first camera pair are performed for the images captured by the second and third sector camera pairs.

In step 1172 the encoded images generated from the captured images are streamed to a playback device along with the information indicating the UV maps to be used in rendering the encoded images being streamed. In some embodiments before a UV map is used it is communicated in the content stream prior to the encoded image for which it is being supplied. Thus in some embodiments rather than being supplied with the UV maps via a separate channel or set of information in some embodiments the UV maps are embedded in the content stream used to deliver the encoded images to requesting playback device or devices.

FIG. 12 illustrates a method 1200 of operating a playback device or system, which can be used in the system of FIG. 1, in accordance with one exemplary embodiment. The method 1200 begins in start step 1202. In step 1204 the playback device transmits a request for content, e.g., to the streaming server of FIG. 1. The playback device then receives in step 1206 various information which may be used for rendering images. For example environmental model information may be received in step 1206 as well as one or more UV maps corresponding to different resolution

25

allocations for one or more regions of the environment. Thus, in step **1206**, the playback device may receive environmental model and/or UV map information corresponding to different resolution allocations. The information received in step **1206** is stored in memory for use on an as needed basis.

Operation proceeds from step **1206** to step **1208** in which one or more images are received, e.g., image captured of an environment to be simulated while an event was ongoing in the environment. In step **1210** information indicating which UV maps are to be used for rendering the one or more received images is indicated. In some embodiments the information indicates which UV map in a set of UV maps corresponding to different resolution allocation which may have been used for a portion of the environment is to be used for rendering left and right eye images of a frame pair, e.g., corresponding to a front portion of an environment. In step **1212** one or more of the received images are decoded.

Operation proceeds from step **1212** to steps **214** in which the decoded images corresponding to surfaces of the environmental model are applied using one or more UV maps corresponding to the indicated resolution allocation that was used to generate the decoded image or images. Operation proceeds from steps **1214** to step **1218** in which image areas corresponding to different portions of the 360 degree simulated environment are combined to the extent needed to generate a contiguous image of a viewing area to be displayed. Then in step **1220** the images are output to a display device with, in the case of stereoscopic image content, different images being displayed to a user's left and right eyes. Operation process from step **1220** back to step **1204** with content being requested, received and processed on an ongoing basis.

FIG. **19** is a drawing **2100** that illustrates mapping of an image portion corresponding to a first sector to the corresponding 120 degree portion of the sphere representing the 3D viewing environment.

In step **1216**, images corresponding to different portions of the 360 degree environment are combined the extent needed to provide a contiguous viewing area to the viewer, e.g., depending on head position. For example, in step **1218** if the viewer is looking at the intersection of two 120 degree sectors portions of the image corresponding to each sector will be seamed and presented together to the viewer based on the know angle and position of each image in the overall 3D environment being simulated.

FIG. **20** is an illustration **2200** showing the result of applying textures to mesh models to form a complete 360 degree view of an environment which may be presented to a user viewing the environment from the perspective of being located in the center of the illustrated environment and with the images applied to the inside of the spherical environment. The result of the simulation and display is a complete world effect in which a user can turn and look in any direction.

The mapped images are output to a display device in step **1220** for viewing by a user. As should be appreciated the images which are displayed will change over time based on the received images and/or because of changes in head position or the user selected viewer position.

Methods and apparatus for supporting delivery, e.g., streaming, of video or other content corresponding to an environment are described. In some embodiments the images corresponding to the environment which are communicated to a playback device exceed the area a user can view at a given time so that content is available in the event the user changes his/her viewing angle by, for example,

26

moving his/her head. By providing images for an environmental area larger than that which can be viewed by a user at a given time the playback device has enough information to provide images should the user's viewing angle change without the playback device having to wait for new images or other content corresponding to a portion of the environment which the user was not previously viewing.

In at least some embodiments the environment is represented using a mesh. Images are captured and encoded into frames, e.g., frames intended for viewing by a left eye and frames intended to be viewed by a right eye. While the techniques are described in the context of 3D stereoscopic applications, the methods can be used for stereoscopic viewing as well with a single stream of frames being communicated rather than a stream of frame pairs.

In some embodiments the techniques are used to communicate images corresponding to a 360 degree viewing area. However, the techniques may be used for communicating images corresponding to less than a 360 degree viewing area, e.g., with a single frame communicating image content corresponding to the 360 degree viewing area. The methods and apparatus of the present invention are particularly well suited for streaming of stereoscopic and/or other image content where data transmission constraints may make delivery of 360 degrees of content difficult to deliver at the maximum supported quality level, e.g., using best quality coding and the highest supported frame rate. However, the methods are not limited to stereoscopic content.

In various embodiments images corresponding to a 360 degree or other area are captured and combined to form an image of the area. The different portions of the image content of the area, e.g., a 360 degree environment, are mapped to a frame which is to be encoded and transmitted. Separate frames may be generated and transmitted for each of the left and right eye views. While the image content corresponding to different portions of the area may have been captured at the same resolution, the mapping of the captured images to the frame may, and in some embodiments is, different for different areas of the environment. For example, the front view portion of the environment may be preserved at full or near full resolution, with the sides and back being incorporated into the frame at lower resolutions. Images corresponding to the top and bottom of a 360 degree environment may be incorporated into the frame at a different, e.g., lower, resolution than the front and/or side views. In some embodiments images corresponding to the top and bottom of an environment are sent separately and, in many cases, as static images or at a different rate than images corresponding to the other portions of the environment.

As a result of the mapping process, a frame communicating an environment may use different numbers of pixels to represent the same size area of a physical environment. For example, a larger number of pixels may be used to represent a forward viewing area with a lower number of pixels being used to represent a rear viewing area. This represents selective downsampling at the time of generate a frame representing the multiple image areas.

In a decoder a the image is mapped or wrapped onto a 3D model of the environment as part of the display process in some embodiments. The map is sometimes referred to as a UV map since UV coordinates are used in mapping the two dimensional frame that is communicated to XYZ space of a 3D model of the environment. The grid (UV map) used to map the transmitted frame to the 3D model takes into consideration the reduced resolution used in some embodiments for the back and side portions of the environment.

In various embodiments, the map used to wrap a communicated frame unto the model of the environment may change to reflect the different allocations of resolution to different portions of the environment. For example, portions of the environment having high motion may be allocated more resolution at points in time when there is high motion and less resolution at other times.

Information on how the transmitted frame should be processed by the decoder to take into consideration the allocation of different amounts of resources, e.g., pixels, to different image areas at different points in time is communicated to the playback device and used to interpret the communicated frame and how it should be applied to the 3D environment.

The method used in various embodiments may be referred to as use of selective resolution allocation in a Panoramic Image map. This approach allows the encoder and playback device to use a UV map to optimize the resolution in a equi rectangular projection so that more of the limited number of pixels available in a communicated frame are used for the more important image element(s) and pixels aren't wasted on image areas of low importance. The methods and apparatus are particularly well suited for devices with limited pixel buffers, such as phones where every pixel is precious because of the phones fairly limited pixel buffer which is available for decoding images.

The process of selective resolution allocation in a panoramic image map can be understood when FIG. 21 is considered. FIG. 21 shows a 3D mesh model **2300** of an environment onto which captured images are to be wrapped as part of the process of rendering an image on a display device. The 3D model **2300** includes a sky mesh **2300**, a 360 degree panoramic mesh **2308** and a ground mesh **2310**. As part of the process of communicating images corresponding to the 3D environment represented by the 3D model, a frame representing an image of the sky is transmitted. A map is used to determine which parts of the transmitted frame are applied to which segments of the sky mesh. In at least one embodiment the sky map includes one segment for each segment of the sky mesh and provides a method of determining which portion of a frame representing an image in what is sometimes referred to as UV space will map to the segments of the sky mesh **2306**. In some embodiments the frame representing the image of the sky is sent once and is thus static or sent at a low rate much less frequently than images to be mapped to the 360 degree panoramic mesh portion of the model **2300**.

As part of the process of communicating images corresponding to the 3D environment represented by the 3D model, a frame representing an image of the ground is transmitted. A ground map is used to determine which parts of the transmitted frame are applied to which segments of the ground mesh. In one embodiment the ground map includes one segment for each segment of the ground mesh **2310** and provides a method of determining which portion of a frame representing an image in what is sometimes referred to as UV space will map to the segments of the ground mesh **2310**. In some embodiments the frame representing the image of the ground is sent once and is thus static or sent at a low rate much less frequently than images to be mapped to the 360 degree panoramic mesh portion of the model **2300**.

Of particular importance are frames corresponding to the 360 degree mesh portion since this includes the areas of the environment which tend to be most frequently viewed. While the image of this environmental area may be captured at a consistent resolution as represented by the uniform

segments of the uncompressed panoramic image map **2302**, different areas to which the panoramic image and panoramic mesh correspond may be of different amounts of importance at different times. For example, frontal areas were the main action is ongoing and/or areas with high motion may be important to represent in detail while other environmental areas may be less important. The uniform allocation of limited resources in terms of pixels of a frame to different areas of an environment is wasteful when the importance of the different image areas is taken into consideration along with the fact that the pixels of the frame are a limited resource. In order to make efficient use of the available pixels of a frame to communicate an image corresponding to a 360 degree environment, a map may be used to allocate different numbers of pixels to different portions of the 360 degree mesh. Thus some portions of the mesh **2308** may be coded using more pixels than other portions. In accordance with one such embodiment, a panoramic image map **2304** with non-uniform segments sizes may be used. While in the case of the FIG. 21 map each segment of the map will be used to map pixels of a received frame to a corresponding segment of the panoramic mesh **2308**, some segments will use more pixels from the transmitted image than other segments. For example, more pixels will be allocated to the middle portions of the panoramic mesh in the FIG. 21 example as represented by the larger segment sizes towards the middle of the map **2304** than towards the top and bottom of the map **2400**. While the map **2304** is used to map portions of a received frame to the mesh **2308**, prior to encoding of the communicated frame one or more segments of the uncompressed image of representing the panoramic environment will be downsampled taking into consideration the panoramic image map. For example, portions of an uncompressed image representing the top and bottom portions of the environment will be downsampled to reflect the small number of pixels allocated in the panoramic image map for representing such image portions while other portions may be subject to lower or no downsampling.

The panoramic image map is generated in some embodiments based on scene analysis and/or taking into consideration a user viewing position. The panoramic image map may be and in some embodiments is changed over time as the location of the main action changes, e.g., ball position during a sporting event in a stadium environment changes. The change is normally limited to occur on a group of pictures boundary within a video stream and/or upon a scene change boundary such as a boundary associated with the start or end of a commercial break in a video sequence. The new map to be used for interpreting frames may be transmitted at a playback device with or prior to a frame which is constructed taking into consideration the new map. Alternatively the playback device may so a variety of predetermined maps which may be used for mapping received frames to the mesh model of the environment and the video stream may include information indicating which of the plurality of maps is to be used for a particular set of communicated frames.

The selective allocation and varying of the image map to take into consideration content and/or user viewing position, can be applied to a full 360 degree area or some small portion of the environment. Accordingly, while shown with an example that maps a frame to a 360 degree environmental area the same method may be applied to a map that is used to map a frame to a 180 degree image area or some other portion of a 360 degree environment. While the map used to map an image to corresponding segments of a 3D model may change, the model need not change. However, changes

29

in the model may be made to reflect changes in the environment, e.g., when a stage is moved and/or other changes in the environment are detected. Thus, both map and model changes are possible.

While different resolution allocations to the top and bottom portions of an area, e.g., the panoramic 360 degree middle area may be made, different resolution allocations may be made within a horizontal area of the environment.

For example, at different times it may be desirable to allocate different amounts of resolution to different portions of a sports field depending on where the ball or area of action is located.

FIG. 21 shows how selective resolution can be used with regard to an image, e.g., frame, which maps to an environmental grid corresponding to, for example, a 360 spherical panoramic mesh. Separate images may be communicated for applying as textures to the sky and ground mesh portions of the world model shown in FIG. 21.

The panoramic image 2302 prior to compression, corresponding to the 360 degree panoramic mesh 2308 includes image content at a generally uniform resolution in the example. In an actual embodiment it should be appreciated that the use of a fisheye lens may introduce some distortions and thus differences in resolution due to lens issues. However, for purposes of explaining the invention it will be presumed that image capture results in an image with a uniform resolution. The grid applied to the panoramic image 2302 is uniform and if used as a UV map would result in uniform resolution allocation to the segments of the 360 degree panoramic portion of the mesh model 2308. However, since a user is less likely to be looking at the bottom or top portions of the environment corresponding to the 360 degree panoramic mesh area, prior to encoding and transmission to the playback device the upper and lower portions are subject to a resolution reduction operation and the UV map to be used during playback is adjusted accordingly. Thus, in mesh 2304 which represents a UV map to be used to render a resolution adjusted image corresponding to the 360 panoramic area of the mesh model, the grid sizes are smaller. Thus, when applied fewer pixels will be extracted for a top segment from the source image and applied to the corresponding segment of the environment than will be extracted and applied for a segment corresponding to the middle horizontal portion of the 360 panoramic mesh model. Thus the UV model takes into consideration the selective allocation of resolution applied to the captured image representing the 360 panoramic area.

The playback device will use the UV mesh which reflects the resolution reduction applied to an image prior to transmission to the playback device when rendering the received image, e.g., applying the received image as a texture, onto the surface of the environmental model, e.g., mesh model of the environment.

While a static UV map reflecting a resolution reduction operation may be and is used in some embodiments, it may be desirable in at least some embodiments where the portion of the environment with the highest priority may change to support the dynamic selection of a resolution allocation approach to use and to use a UV map corresponding to the selected resolution allocation. In such a way, resolution allocation may be changed to reflect which portion of the environment is given priority in terms of resolution at a given time.

FIG. 22 represented by reference number 2400, shows a first captured image 2402 of a first portion of an environment. Each large dot represents a pixel. The image 2402 is of uniform resolution as represented by the 4 pixels in each

30

square grid area. Small dots are used to indicate that the image continues and extends toward the other illustrated portions of the image 2402. When a first resolution allocation is selected, e.g., a resolution which gives priority to the middle portion of the image 2402, resolution will be preserved during the middle portion of the image 2402 but reduced for the left and right portions. Such a resolution allocation may be desirable where, for example, the image 2402 is of a sports field and the action is at the center portion of the sports field when image 2402 is captured. Arrows extending from image 2402 towards reduced resolution image 2404 represent the application of a first selective resolution reduction operation to image 2402 to generate image 2404. The first resolution reduction operation may involve a downsampling applied to the left and right portions of image 2402 but not the middle portion. The grid shown as being applied to image 2404 represents the resolution allocation used to generate image 2404 from image 2402. As can be seen the first resolution adjusted image 2404 includes half as many pixel in the two left and right most rows of the image as did image 2402 but the same number of pixels for segments towards the center portion of the image 2404. Grid 2406 represents a first UV map corresponding to the first resolution allocation which is suitable for mapping segments of the image 2404 to segments of the model of the environment.

FIG. 23 represented by reference number 2500, shows a first captured image 2502 of the first portion of the environment. As in the case of FIG. 22, each large dot represents a pixel. The image 2502 is of uniform resolution as represented by the 4 pixels in each square grid area. Small dots are used to indicate that the image continues and extends toward the other illustrated portions of the image 2502. When a second resolution allocation is selected, e.g., a resolution which gives priority to the left and middle portions of the image 2502, resolution will be preserved in the left and middle portions of the image 2502 but reduced for the right portion. Such a resolution allocation may be desirable where, for example, the image 2502 is of a sports field and the action is at the left portion of the sports field when image 2502 is captured. Arrows extending from image 2502 towards reduced resolution image 2504 represent the application of a second selective resolution reduction operation to image 2502 to generate image 2504. The second resolution reduction operation may involve a downsampling applied to the right portion of image 2502 but not the left or middle portions. Note that the area to which the downsampling is applied is of a size equal to the area to which downsampling was applied in FIG. 22. As a result images 2404 and 2504 will have the same total number of pixels but with the resolution being different in different areas of the images 2404, 2504.

While total pixel count is maintained as being constant for different reduced resolution images with the resolution allocation applying to different areas of an image, this is not critical and different images may include different numbers of pixels after a resolution reduction operation. However, keeping the pixel count constant facilitates encoding since the encoder can treat the images to be encoded as being of the same size even though when used in playback device different portions of the model will be rendered at different resolutions due to the use of different UV maps for different resolution allocations.

The grid shown as being applied to image 2504 represents the resolution allocation used to generate image 2504 from image 2502. As can be seen the second resolution adjusted image 2504 includes half as many pixel in the four right

31

most rows of the image as did image **2502** but the same number of pixels for segments towards the left and center portions.

Grid **2506** represents a first UV map corresponding to the first resolution allocation which is suitable for mapping segments of the image **2504** to segments of the model of the environment.

FIG. **24** represented by reference number **2600**, shows a first captured image **2602** of the first portion of the environment. As in the case of FIGS. **22** and **23**, each large dot represents a pixel. The image **2602** is of uniform resolution as represented by the 4 pixels in each square grid area. Small dots are used to indicate that the image continues and extends toward the other illustrated portions of the image **2602**. When a third resolution allocation is selected, e.g., a resolution which gives priority to the middle and right portions of the image **2602**, resolution will be preserved in the middle and right portions of the image **2602** but reduced for the left portion. Such a resolution allocation may be desirable where, for example, the image **2602** is of a sports field and the action is at the right portion of the sports field when image **2602** is captured. Arrows extending from image **2602** towards reduced resolution image **2604** represent the application of a third selective resolution reduction operation to image to generate image **2604**. The third resolution reduction operation may involve a downsampling applied to the left portion of image **2602** but not the right or middle portions. Note that the area to which the downsampling is applied is of a size equal to the area to which downsampling was applied in FIGS. **22** and **24**. As a result images **2604** will have the same total number of pixels as images **2404**, **2504** but with the resolution being allocated differently in terms of the portion of the environment to which higher resolution is allocated.

The grid shown as being applied to image **2604** represents the resolution allocation used to generate image **2604** from image **2602**. As can be seen the third resolution adjusted image **2604** includes half as many pixel in the four left most rows of the image as did image **2602** but the same number of pixels for segments towards the right and center portions.

Grid **2606** represents a first UV map corresponding to the first resolution allocation which is suitable for mapping segments of the image **2604** to segments of the model of the environment.

UV map **2406** is communicated to a playback device for use with an image generated using the first resolution allocation. UV map **2406** is communicated to a playback device for use in rendering an image generated using the second resolution allocation and UV map **2606** is communicated to the playback device for use in rendering an image generated using the third resolution allocation. The streaming system and the playback system both store the set of UV maps **2406**, **2506**, **2606** with the streaming system indicating which UV map should be applied to which image and the rendering device, e.g., playback device, using the indicated UV map associated with a received image.

While different resolution allocation are supported through the use of different UV maps this can be transparent to the decoder in the playback device which decodes received images since the decoder need not have knowledge of which of the plurality of possible resolution allocations were used to generate a received encoded image which is to be decoded by the decoder in the playback device.

FIG. **25** which comprises FIGS. **25A** and **25B** illustrates an exemplary method **2900** of operating a content processing and delivery system in accordance with an exemplary embodiment. FIG. **25A** shows the first part of method **2900**.

32

FIG. **25B** shows the second part of method **2900**. The method **2900** shown in FIG. **25** starts in step **2902** with the content processing and delivery system being initialized to process and deliver content, e.g., image content and/or information used to render images. In some embodiments the method of flowchart **2900** is performed using the content delivery and processing system **700** of FIG. **7**.

From start step **2902** operation proceeds to steps **2904** and **2906**, which may be performed in parallel and, optionally, asynchronously. In various embodiments customer rendering and playback devices are provided with information that can be used in rendering of image content and/or providing 3D playback experience to the viewers. In some embodiments this includes providing environmental model and/or other environmental information to the customer devices to be used in rendering images as part of an environmental simulation. In step **2904** a 3D environmental model and/or information that can be used to model is communicated to one or more customer devices. In some embodiments the model is a mesh model of the environment from which one or more images are captured. In some embodiments additional information which can be used in rendering images, e.g., one or more UV maps are also communicated to the customer devices, e.g., content playback devices, in step **2905**. The UV maps correspond to different resolution allocations with different UV maps, also referred to as texture maps, providing different mappings of pixels of transmitted images to segments of the environmental model. If the UV maps are communicated in step **2905** they can later be identified when they are to be used to map a transmitted image and need not be retransmitted multiple time to the playback device. However, in some embodiments a set of UV maps is not communicated in step **2905** and an applicable UV map is transmitted with or prior to communication of an image to which the UV map is to be applied and used.

In some embodiments the information in steps **2904** and **2905** is communicated once, e.g., prior to communicating actual image content to the customer devices. While environmental map information and/or environmental models may be communicated to the playback device in some embodiments where such information is generated and/or available at the server side, in some other embodiments the environment may be assumed to be a default size and shape, e.g., a sphere or half sphere and in such a case the default environmental module and/or UV maps may be preloaded in the playback device and need not be transmitted by the server.

The processing of image content begins in step **2906** which can be performed in parallel with steps **2904**, **2905**. In step **2906** image content is received by the processing system, e.g., content delivery system **700** shown in FIG. **7**. The image content received in step **2906** may be from an image capturing device such as the ones discussed in the application such as the one shown in FIG. **13**. In some embodiments the step **2906** of receiving image content includes step **2908** wherein a first image corresponding to a first portion of an environment, e.g., environment of interest where images are captured, is received. In some embodiments the first image is one image of an image pair that also includes a second image, with the first image being one of a left and right eye image pair, the second image being a second one of a left and right eye image pair. In some embodiments the first and second images are received as part of the image pair in step **2906**. Thus in some such embodiments step **2906** further includes step **2910** of receiving the second image.

33

Operation proceeds from step **2906** to step **2912** the system selects a first resolution allocation to be used for at least one image corresponding to a first portion of the environment. This selection may be and sometimes is based on detection of motion in the received image content, the location of particular objects such as a sports jersey, and/or human input indicating which portion of the captured image is to be given priority and preserved at a higher resolution during encoding. For example, detection of player's jerseys or uniforms may indicate areas to be preserved at high resolution in which case a resolution allocation which preserves the areas where the uniforms are detected may and in some embodiments will be selected. Other portions of the image may be and sometimes are subject to down sampling. Each resolution may correspond to a particular UV map which is intended to be used for mapping images produced by using a particular corresponding resolution allocation.

Operation proceeds from step **2912** to step **2914**. In step **2914** in which it is determined if the selected first resolution allocation is different from a previously selected resolution allocation, e.g., indicative of a change in down sampling and UV map. The selected first resolution allocation may be one of a plurality of supported resolution allocations corresponding to different supported UV maps corresponding to the first portion of the environment captured in the first image. In accordance with one aspect from the plurality of supported resolution allocations a resolution allocation may be selected at a given time to process a current image and/or group of images. If it is determined that the selected first resolution allocation is different than the previously selected resolution allocation the operation proceeds from step **2914** to step **2916** where new downsampling and/or filtering information corresponding to the newly selected resolution allocation used to control resolution reduction is loaded and then operation proceeds to step **2918**. If in step **2914** it is determined that the selected first resolution allocation is the same as the previously selected resolution allocation (or is the same as a default allocation if no previous selection was made) then there is no need for new down sampling and/or filtering information to be loaded and thus the operation proceeds directly to step **2918**. The selected resolution allocation for an image indicates how down sampling is to be applied to an image which is to be encoded and transmitted to the playback device.

In step **2918** a resolution reduction operation, e.g., down-sampling, is performed on the first image of the first portion of the environment in accordance with the selected first resolution allocation to generate a first reduced resolution image **2919**. The first reduced resolution image **2919** which is the output of step **2918** includes at least some image portions having different resolutions.

Operation proceeds from step **2916** to step **2920** in embodiments where pairs of images are processed, e.g., stereoscopic image pairs including left and right eye views. In step **2920** a resolution reduction operation is performed on the second image of the first portion of the environment, e.g., the second image in stereoscopic image pair, in accordance with the selected first resolution allocation to generate a second reduced resolution image **2921**. The second reduced resolution image **2921** which is the output of step **2918** includes at least some image portions having different resolutions. Thus, where stereoscopic image pairs are processed, both the left and right eye images of a pair will be subject to the same resolution reduction operation.

While step **2920** is shown as being performed after step **2918** it may be performed in parallel with step **2918** simultaneously. The data output of steps **2918** and **2920**, e.g., the

34

generated first and second reduced resolution images **2919** and **2921**, serve as inputs to the next step **2922**. In the case of non-stereo image content, a single image will be processed and the second image will not be present.

In step **2922** the reduced resolution image **2919** and/or reduced resolution image **2921** are encoded. In step **2924** the first reduced resolution image is encoded. In step **2926** the second reduced resolution image, when present, is encoded.

Operation proceeds from step **2922** to step **2928**. In step **2928** the encoded reduced resolution images are stored in memory, e.g., for subsequent communication, e.g., streaming to a content playback device, e.g., located at a customer premises such as a house or home. Operation proceeds from step **2928** to step **2930** via connecting node B **2929**. In step **2930** the encoded reduced resolution image(s) are communicated to a playback device. This may involve transmitting, e.g., streaming, the images to the playback device over a wired network, cable network or wireless network or some other type of network. Step **2930** includes steps **2932** and step **2934**. In step **2932** the first reduced resolution image is communicated to the customer playback device, e.g., in encoded form and in step **2934** in the second reduced resolution image is communicated to the playback device, e.g., in encoded form. Step **2934** is performed when a stereo pair of images is communicated, e.g., in a single frame or pair of frames.

Operation is shown proceeding from step **2930** to step **2936**. However depending on the embodiment step **2936** may precede step **2930**. In step **2936** a texture map, e.g., first texture map, to be used to map the encoded images to the model of the environment is indicated or provided to the playback device. The identification of the first texture map may be sufficient where the first texture map, e.g., UV map, was already loaded into the playback device e.g., as part of step **2905**. Based on the communicated information and/or map, the playback device knows that it is to use the first UV map with the first and second images which were produced using the first resolution allocation to which the first UV map corresponds. The first UV map may be used by the playback device to render other images which are also produced in accordance with the first resolution allocation. In some embodiments a resolution allocation is maintained for a group of pictures and thus the same UV map may be used for multiple consecutive images in such embodiments.

Operation proceeds from step **2936** in which the playback device is provided information about what texture map to use while rendering the first and second images to step **2938** which relates to processing of an additional image or images, e.g., a third image and/or fourth image. The third and/or fourth image may be and in some embodiments are left and right images of a second stereoscopic image pair or some other image or images of the environment captured after the first image.

In step **2940** a second resolution allocation is selected to be used for the received images, e.g., third and/or fourth images. The resolution allocation may be determined using the same techniques used to determine the first resolution allocation, e.g., identifying an area or areas of importance based on motion, presence of an object such as sports jersey, ball, etc. Once the second resolution allocation is selected from the set of resolution allocations, e.g., each corresponding to a different UV map, operation proceeds to step **2942**. In step **2942** a check is made to determine if the second resolution allocation is different from the first resolution allocation. The second resolution allocation may be different, e.g., because the ball or players may have moved to a different portion of the field since the first image was

captured. If the second selected resolution allocation is different than the first selected resolution allocation new downsampling information needs to be loaded and used and operation proceeds to step 2944. In step 2944 the new downsampling and/or other resolution allocation information is loaded so that it can be used in the resolution reduction step 2946. If in step 2942 it is determined that the second resolution allocation is the same as the first, the processor of the system implementing the method 2900 already knows the downsampling to be performed since it was used process the first image and need not load new downsampling information and operation proceeds to step 2946.

In step 2946 a resolution reduction operation, e.g., downsampling, is performed on the received third and/or fourth image to produce reduced resolution versions of the third and/or fourth images 2947. Operation proceeds from step 2946 to step 2948 in which the reduced resolution third and/or fourth images are encoded prior to being communicated, e.g., transmitted, to the playback device in step 2950.

In step 2952, which is shown being performed after step 2950 but which may and sometimes does precede step 2950 or occur in parallel with step 2950, the information indicating the UV map to be used for rendering the third and fourth images is communicated to the playback device. This may involve sending the UV map to be used to the playback device or simply identifying a previously stored UV map. Since the third and fourth images were generated using the second resolution allocation the information will identify the UV map corresponding to the second UV allocation. Operation proceeds from step 2952 via connecting node 2954 to step 2906 where additional image content is received, e.g., from a camera device, and treated as new first and second images.

Over time a sequence of images representing view may be received and processed with the resolution allocation used at a given time depending on the received image content and/or user input. Over time as different resolution allocations are used, the content playback device will be signaled to use different corresponding UV maps. Thus when the second resolution allocation is different from the first resolution allocation the playback device will be instructed to use a second different UV map to render images generated in accordance with the second resolution allocation which is different from a first UV map used to render images generated in accordance with the first resolution allocation. A large number of different resolution allocations can be used in combination with corresponding UV maps allowing for a wide variety of different resolution allocations to be supported.

FIG. 26 illustrates an exemplary embodiment of a content playback method 2700 which may be, and in some embodiments is, implemented on exemplary computer system/content playback device 800. The method 2700 may be used by a device which receives content encoded and transmitted in accordance with the method of FIG. 25.

For explanatory purposes, the exemplary content playback method 2700 will be explained in connection with the playback device 800 shown in FIG. 8. It should be appreciated that the method 2700 can be implemented on other apparatus. The exemplary playback method 2700 begins in start step 2702 from which operation proceeds to step 2704.

In step 2704, the receiver 860 of the network interface 810 of the content playback device 800 receives a mesh model of an environment. Operation proceeds from step 2704 to step 2706. In step 2706, the receiver 860 of the network interface 810 of the content playback device 800 receives

one or more image maps, e.g., one or more image UV maps, indicating a mapping between an image and the mesh model of an environment. In some embodiments, step 2706 includes sub-step 2708 and/or sub-step 2710. In sub-step 2708, the receiver 860 of the network interface 810 of the content playback device 800 receives a first image map. In sub-step 2710, the receiver 860 of the network interface 810 of the content playback device 800 receives a second image map. Operation proceeds from step 2706 to step 2712.

In step 2712, the content playback device 800 stores the received image map or maps in a storage device, e.g., memory 812. For example, UV MAP 1 836 and UV MAP 2 836 are stored in memory 812. In some embodiments the received image maps are stored in a storage device coupled to the content playback device 800. Operation proceeds from step 2712 to step 2714.

In step 2714, the receiver 860 of the network interface 810 receives an encoded image. Operation proceeds from step 2714 to step 2716. In step 2716, the decoder 864 of the playback device 800, decodes the received encoded image. In some embodiments, a hardware decoder module decodes the received encoded images. In some embodiments, the processor 808 executing instructions from decoder module 820 decodes the received encoded image. Operation proceeds from step 2716 to step 2718.

In step 2718, the decoded image is mapped to the mesh model of the environment in accordance with the first image map to produce a first rendered image. The first image map mapping different numbers of pixels of the decoded image to different segments of the mesh model of the environment. While the mapping of the different numbers of pixels of the decoded image to different segments of the mesh model of the environment may occur in a variety of different ways, in some embodiments, the different numbers of pixels are mapped to environmental regions of the same size but located at different locations in the environment. In some such embodiments, segments in the environment corresponding to action are allocated more pixels than segments in which less or no action is detected. In some embodiments, at least some segments corresponding to a front viewing area are allocated more pixels per segment than segments corresponding to a rear viewing area. This mapping may be, and in some embodiments is, performed by the processor 808 of the playback device 800. Operation proceeds from step 2718 to step 2719.

In step 2719, the first rendered image is displayed for example on display 802 of content display device 800.

In some embodiments, operation proceeds from step 2719 to step 2720. In step 2720, the receiver 860 of the network device 810 of the playback device 800 receives a signal indicating that a second image map should be used to map portions of received frames to the environmental mesh model. In some embodiments the decoded image is a frame. Operation proceeds from step 2720 to optional step 2722. In step 2722, in response to receiving the signal indicating that a second image map should be used to map portions of received frames to the environmental mesh model, a second image map is used to map portions of received frames to the environmental mesh model to produce one or more additional rendered images, e.g., a second rendered image. In some embodiments, the second image map is the second image map received in step 2710.

In some embodiments, the first image map allocates a first number of pixels of a frame to a first segment of said environmental mesh model wherein the decoded image is a frame and said second image map allocates a second number of pixels of the frame to the first segment of said environ-

37

mental mesh model, the first and second number of pixels being different. The mapping of step 2722 may be, and in some embodiments is, performed by the processor 808 of the playback device 800. Operation proceeds from optional step 2722 to optional step 2724.

In step 2724, the additional rendered image(s), e.g., the second rendered image, is displayed for example on display 802 of content display device 800. Operation proceeds from step 2724 to step 2704 where the method continues as previously described.

In some embodiments of the exemplary method 2700, the received mesh model of an environment is stored in a storage device, e.g., 3D environmental mesh model 832 stored in memory 812 of the playback device 800. In some embodiments, the received encoded image data which may be, and in some embodiments are encoded scene portions, is stored in a storage device, e.g., received encoded data 824 stored in memory 812 of the playback device 800. In some embodiments, the decoded image data is stored in a storage device, e.g., decoded data 826 stored in memory 812 of the playback device 800. In some embodiments, the one or more of the rendered images are stored in a storage device, e.g., memory 812 of the playback device 800. In some embodiments, the first and second images are rendered by the processor 808 executing instructions contained in the image generation module 822. In some embodiments, a hardware, e.g., circuits, image generation module performs the operation of rendering the one or more images, e.g., the first and/or second rendered images.

The exemplary embodiment of method 2800 of communicating information to be used to represent an environment will now be described in connection with FIG. 10. The exemplary method 2800 may be, and in some embodiments is, implemented by a content delivery system such as for example content delivery system 700 illustrated in FIG. 7.

Operation of the method 2800 begins in start step 2802. Operation proceeds from step 2802 to step 2804.

In step 2804, a first image map to be used to map portions of a frame to segments of an environmental model are communicated, e.g., to a content playback device such as for example content playback device 800 illustrated in FIG. 8. The first image map allocates different size portions of the frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of the environmental model. In some embodiments, the network interface 710 of the content delivery system 700 performs this operation. In such embodiments, the network interface 710 includes a transmitter 711 which performs this function. Operation proceeds from step 2804 to step 2806.

In step 2806, a first frame including at least a portion of a first image to be mapped to the environmental model using the first image map is communicated, e.g., to the content playback device 800. In some embodiments, the network interface 710 of the content delivery system 700 performs this operation. In some embodiments, the network interface 710 includes a transmitter 711 which performs this operation. Operation proceeds from step 2806 to step 2808.

In step 2808, a second image map to be used to map portions of a frame to segments of the environmental model is communicated, e.g., to the content playback device such as for example content playback device 800. The second image map allocates different size portions of the frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said model. The second image map allocates a different number of pixels to a first segment of the environmental

38

model than are allocated by the first image map. In some embodiments, the network interface 710 of the content delivery system performs this operation. In some embodiments, the network interface 710 includes a transmitter 711 which performs this operation. Operation proceeds from step 2808 to step 2810.

In step 2810, a second frame including at least a portion of a second image to be mapped to the environmental model using the second image map is communicated e.g., to the content playback device such as for example content playback device 800. The first and second image maps map different numbers of pixels to an area corresponding to the same portion of an environment thereby providing different resolution allocations for the same portion of the environment based on which of the first and second image maps are used. In some embodiments, the network interface 710 of the content delivery system performs this operation. In some embodiments, the network interface 710 includes a transmitter 711 which performs this operation. Operation proceeds from step 2810 to step 2804 where operation proceeds as previously described.

FIGS. 27, 28 and 29 show how a playback device, such as the playback device or devices shown in any of the other figures, can perform image rendering using a UV map corresponding to the resolution allocation that was used to generate the image to be rendered.

FIGS. 27 shows how a reduced resolution image 2404 can be rendered using the UV map 2406 and an environmental module 3002 with environmental segments in the model corresponding to segments of the UV map. The top portion of FIG. 27 shows the relationship between segments of the UV map 2406 and the segments of the environmental model 3002. A first segment of the UV map 2406 corresponds to a first environmental module segment (EMS 1) of environmental model 3002, as represented by the solid arrow extending from the first segment of the UV map 2406 and EMS 1. A second environmental module segment (EMS 2) of environmental model 3002 corresponds to the second segment of the UV map 2406 as indicated by the dashed arrow extending from the second segment of the UV map 2406 and EMS 2. A third environmental module segment (EMS 3) of environmental model 3002 corresponds to the third segment of the UV map 2406 as indicated as represented by the dashed arrow extending from the second segment of the UV map 2406 and EMS 3. There is a known, e.g., one to one, relationship between other segments of the UV map 2406 and the environmental model 3002.

During rendering, the UV map 2406 is used to determine how to apply portions of an image generated in accordance with the first resolution allocation to portions of the environmental model 3002, as a texture. In the FIG. 27 UV map 2404 is applied to the communicated image 2404 to determine how to segment the image 2404 into sets of pixels to be applied to the corresponding segments of the environmental model 3002. The pixels in the segments of the image 2404 corresponding to a segment of the UV map 2406 are then applied to the corresponding segment of the environmental model 3002, e.g., as a texture, with scaling and reshaping being used as necessary to cover the surface of the segment of the environmental model 3002. The portion of the image applied to the corresponding segment of the environmental model 3002 is scaled and/or adjusted in shape as necessary to fully occupy the corresponding segment of the environmental model 3002 in some embodiments. Thus, for example, two pixels of the communicated image corresponding to the first segment of the UV map 2406 are scaled to fully occupy the first segment EMS1 of the environmental

model **3002** to which they are applied. Similarly in the FIG. **27** example, the two pixels of the image **2404** being rendered, corresponding to the second segment of the UV map **2406** are scaled to fully occupy the second segment EMS2 of the environmental model **3002** to which they are applied as a texture. In the FIG. **27** example, the third segment of the UV map corresponds to four pixels of the image **2404** to be rendered. The four pixels are applied as a texture to the third segment EMS3 of the environmental model **3002** as a texture during the rendering process. Thus, assuming the third segment of the environmental model **3002** is the same size as the first and second segments of the environmental model, the third segment will be of higher resolution than the first and second segments and correspond to more pixels in the received image **2404** than either of the first and second segments. Thus the segments of the UV map **2406** corresponding to portions of an image which were subject to resolution reduction prior to encoding may correspond to the same size area of the environmental model **3002** of another segment which does not correspond to a resolution reduction operation. As should be appreciated the segment corresponding to the area where resolution reduction was not performed will be displayed in the generated image of the simulated environment at a higher resolution than the portion to which resolution reduction was performed prior to encoding.

As discussed above, different resolution reduction operations may be performed to produce images that are transmitted. The playback device will use a UV map corresponding to the resolution reduction operation that was performed when rendering the received images. Thus, while the environmental model **3002** may remain the same for multiple images, different UV maps **2406**, **2506**, **2606** may be used with the same environmental model **3002**.

FIG. **28** shows the application of UV map **2506** to an image **2504**, generated using the second selective resolution reduction operation, which allocates less resolution to the right portion of an image corresponding to a portion of an environment than the left and middle portions. Thus FIG. **28** shows how a reduced resolution image **2504** can be rendered using the UV map **2506** and the environmental model **3002** with environmental segments corresponding to segments of the UV map. The top portion of FIG. **28** shows the relationship between segments of the UV map **2506** and the segments of the environmental model **3002**. A first segment of the UV map **2506** corresponds to the first environmental module segment (EMS 1) of environmental model **3002**, as represented by the solid arrow extending from the first segment of the UV map **2506** and EMS 1. A second environmental module segment (EMS 2) of environmental model **3002** corresponds to the second segment of the UV map **2506** as indicated by the dashed arrow extending from the second segment of the UV map **2506** and EMS 2. A third environmental module segment (EMS 3) of environmental model **3002** corresponds to the third segment of the UV map **2506** as indicated by the dashed arrow extending from the second segment of the UV map **2506** and EMS 3.

During rendering, the UV map **2506** is used to determine how to apply an image to be rendered to the environmental model **3002**. FIG. **28** shows how the communicated image **2504** and the pixels in the segments of the image corresponding to a segment of the UV map are applied to the corresponding segment of the environmental model **3002**. The portion of the image **2504** applied to the corresponding segment of the UV map is scaled and/or adjusted in shape as necessary to fully occupy the corresponding segment of the UV map. Thus, for example, four pixels of the communi-

cated image corresponding to the first segment of the UV map **2504** are scaled to fully occupy the first segment EMS1 of the environmental model to which they are applied. Similarly in the FIG. **28** example, the four pixels of the image being rendered, corresponding to the second segment of the UV map are scaled to fully occupy the second segment EMS2 of the environmental model **3002** to which they are applied as a texture. In the FIG. **28** example, the third segment of the UV map also corresponds to four pixels of the image to be rendered. The four pixels are applied as a texture to the third segment of the environmental model as a texture during the rendering process. Thus, assuming the third segment of the environmental model is the same size as the first and second segments of the environmental model, the third segment will be of the same resolution as the first and second segments. In accordance with the second resolution allocation scheme resolution reduction is not applied to the left and middle portions of the image but resolution reduction is performed with regard to the right side of the image. Thus while the first, second and third segments of the rendered image will be of the same resolution in the FIG. **28** example, segments corresponding to the right side of the image and thus the right side of the environmental model **3002** will be of lower resolution.

FIG. **29** shows the application of UV map **2606** to an image **2604**, generated using the third selective resolution reduction operation, which allocates less resolution to the left portion of an image corresponding to a portion of an environment than the middle and right portions. Thus FIG. **28** shows how a reduced resolution image **2604** can be rendered using the UV map **2606** and the environmental model **3002** with environmental segments corresponding to segments of the UV map **2606**. The top portion of FIG. **29** shows the relationship between segments of the UV map **2606** and the segments of the environmental model **3002**. A first segment of the UV map **2606** corresponds to the first environmental module segment (EMS 1) of environmental model **3002**, as represented by the solid arrow extending from the first segment of the UV map **2606** and EMS 1. A second environmental model segment (EMS 2) of environmental model **3002** corresponds to the second segment of the UV map **2606** as indicated by the dashed arrow extending from the second segment of the UV map **2606** and EMS 2. A third environmental module segment (EMS 3) of environmental model **3002** corresponds to the third segment of the UV map **2606** as indicated by the dashed arrow extending from the second segment of the UV map **2606** and EMS 3.

During rendering, the UV map **2606** is used to determine how to apply an image to be rendered to the environmental model **3002**. FIG. **29** shows how the communicated image **2604** and the pixels in the segments of the image corresponding to a segments of the UV map are applied to the corresponding segments of the environmental model **3002**. The portion of the image **2604** corresponding to a segment of the environmental model **3002** as indicated by the UV map **2606** is scaled and/or adjusted in shape as necessary to fully occupy the corresponding segment of the environmental model **3002**. Thus, for example, two pixels of the communicated image **2604** corresponding to the first segment of the UV map **2606** are scaled to fully occupy the first segment EMS1 of the environmental model to which they are applied. Similarly in the FIG. **29** example, the two pixels of the image being rendered, corresponding to the second segment of the UV map **2606** are scaled to fully occupy the second segment EMS2 of the environmental model **3002** to which they are applied as a texture. In the FIG. **29** example,

41

the third segment of the UV map also corresponds to two pixels of the image to be rendered. The two pixels are applied as a texture to the third segment of the environmental model **3002** as a texture during the rendering process. Thus, assuming the third segment of the environmental model **3002** is the same size as the first and second segments of the environmental model **3002**, the third segment will be of the same resolution as the first and second segments. In accordance with the third resolution allocation scheme resolution reduction is not applied to the middle and right portions of the transmitted image but resolution reduction is performed with regard to the left side of the image. Thus while the first, second and third segments of the rendered image will be of the same resolution in the FIG. **29** example, segments corresponding to the middle and side of the image and right side of the environmental model **3002** will be of higher resolution.

Thus, by using different UV maps different resolution allocation can be achieved during playback while the size and/or number of pixels in the input images remains the same. This provides an easy and efficient way of changing resolution allocations without having to alter the size or number of pixels in the images being transmitted.

Another exemplary apparatus for playing back content will now be described. The apparatus includes a receiver for receiving signals, a mesh model of an environment, one or more image maps, e.g., UV map(s), indicating a mapping between an image and the mesh model of an environment, and one or more encoded images. In some embodiments, the receiver of the apparatus is configured to receive a mesh model of an environment, a first image map, a second image map, and an encoded image. The apparatus also includes or is coupled to a storage device such as a memory for storing received signals, mesh models, image maps, and images such as encoded, decoded and produced images. The apparatus further includes a decoder for decoding received encoded images and a processor configured to map a decoded image to a mesh model of an environment in accordance with a first image map to produce a first rendered image. The first image map mapping different numbers of pixels of the decoded image to different segments of said mesh model of the environment. In some embodiments, the apparatus is configured so that the different numbers of pixels are mapped to environmental regions of the same size but located at different locations in the environment. In some embodiments, the segments in the environment corresponding to action are allocated more pixels than segments in which less or no action is detected. In some embodiments, the apparatus is configured so that at least some segments corresponding to a front viewing area are allocated more pixels per segment than segments corresponding to a rear viewing area. In some embodiments, the apparatus includes or is coupled to a display device on which images produced by the apparatus are displayed. The processor of the apparatus may be, and typically is, configured to operate the apparatus to store received signals, mesh models, image maps, and images such as encoded, decoded and produced images in a storage device included in or coupled to the apparatus.

In some embodiments, the receiver of the apparatus is configured to receive a signal indicating that a second image map should be used to map portions of received frames to said environmental mesh model. The processor of the apparatus is further configured to operate the apparatus in response to the received signal indicating that a second image map should be used to map portions of received frames to the environmental mesh model to use a second

42

image map, typically the second received image map, to map portions of received frames to the environmental mesh model to produce a second rendered image. In some of such apparatus, the decoded image is a frame and the first image map allocates a first number of pixels of the frame to a first segment of the environmental mesh model and the second image map allocates a second number of pixels of the frame to the first segment of the environmental mesh model, the first and second number of pixels being different. The processor the apparatus is typically configured to display the second rendered image to a display which may be either included as part of the apparatus or coupled to the apparatus.

An exemplary apparatus for communicating information to be used to represent an environment will now be discussed. The exemplary apparatus includes a processor configured to operate said apparatus to: (i) communicate a first image map to be used to map portions of a frame to segments of an environmental model, the first image map allocating different size portions of the frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of the model, and (ii) communicate a first frame including at least a portion of a first image to be mapped to said environmental model using said first image map.

In some embodiments, the processor of the apparatus is further configured to operate the apparatus to: (i) communicate a second image map to be used to map portions of a frame to segments of the environmental model, said second image map allocating different size portions of the frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said model, the second image map allocating a different number of pixels to a first segment of said model than are allocated by said first image map, e.g., UV map, and (ii) communicate a second frame including at least a portion of a second image to be mapped to said environmental model using the second image map. In some embodiments of the apparatus, the first and second image maps map different numbers of pixels to an area corresponding to the same portion of an environment thereby providing different resolution allocations for said same portion of the environment based on which of the first and second image maps are used.

In some embodiments, the apparatus is a server providing a real time content stream. In some embodiments, the apparatus is a real time content delivery system including an environmental mesh generation module, a map generation module, e.g., UV map generation module, and an I/O interface and/or a network interface for communicating information including signals, models, maps and images. In some embodiments, the modules include software instructions which when executed cause the processor to perform various routines. In some embodiments, the modules are hardware modules, e.g., circuitry. In some embodiments, the modules are a combination of hardware and software modules.

An exemplary content processing and delivery system, e.g., system **700**, implemented in accordance with one exemplary embodiment comprises: a processor (e.g., processor **708**) configured to: i) select a first resolution allocation to be used for at least one image corresponding to a first portion of an environment; and ii) perform a resolution reduction operation on a first image of the first portion of the environment in accordance with the selected first resolution allocation to generate a first reduced resolution image; and a transmitter (e.g., a transmitter **713** of interface **710**) configured to communicate the first reduced resolution image to a playback device.

In some embodiments selection of a resolution allocation is performed based on a region of importance in the first portion of the environment. In some embodiments the region of importance corresponds to an area of motion in the first portion of the environment. In some embodiments the region of importance is a region indicated by a system operator. In some embodiments the region of importance is a region determined by detecting which portion of the environment included in the first image one or more individuals is looking at prior to or at the time the first image is captured.

In some embodiments the transmitter is further configured to: communicate to the playback device a first texture map (UV map) to be used to map portions of the images generated in accordance with the first resolution allocation to a surface of a model of the environment. In some embodiments the size of a first segment in the first texture map is a function of the amount of resolution reduction applied to a corresponding first area of the first image to generate a first segment of the first reduced resolution image. In some embodiments the first texture map includes a second segment corresponding to a portion of the first image which was not subject to a resolution reduction operation, the size of the second segment in the second texture map being the same as the size of the segment in the first image.

In some embodiments the size of the first segment in the texture map is reduced from the size of the source of the corresponding area in the first image by an amount which is based on the amount of resolution reduction applied to the corresponding first area of the first image. In some embodiments the transmitter is further configured to communicate to the playback device an environmental model. In some embodiments the first texture map corresponds to a portion of the environmental model, the first texture map providing information indicating how to map portions of images subject to the first resolution allocation to a portion of the environmental model. In some embodiments the first image is one image of an image pair the first image and a second image, the first image being one of a left and right eye image pair, the second image being a second one of a left and right eye image pair. In some embodiments the processor is further configured to perform a resolution reduction operation on the second image in accordance with the selected first resolution allocation to generate a second reduced resolution image, and the transmitter is further configured to communicate the second reduced resolution image to the playback device as part of a first stereoscopic image pair.

In some embodiments the processor is further configured to: select a second resolution allocation to be used for another image corresponding to a first portion of the environment, the another image being a third image; and perform a resolution reduction operation on the third image in accordance with the selected second resolution allocation to generate a third reduced resolution image. In some embodiments the transmitter is further configured to communicate the third reduced resolution image to a playback device.

In some embodiments the transmitter is further configured to communicate to the playback device a second texture map (UV map) to be used to map portions of images generated in accordance with the second resolution allocation to the surface of the model of the environment. In some embodiments the size of a first segment in the second texture map is a function of the amount of resolution reduction applied to a corresponding first area of the third image to generate a first segment of the third reduced resolution image. In some embodiments the second texture map includes a third segment corresponding to a portion of the third image which was not subject to a resolution reduction operation, the size

of the third segment in the second texture map being the same as the size of the segment in the third image.

In some embodiments the size of the first segment in the second texture map is reduced from the size of the source of the corresponding area in the third image by an amount which is based on the amount of resolution reduction applied to the corresponding first area of the third image. In some embodiments the second texture map corresponds to the same portion of the environmental model as the first texture map, the second texture map providing information indicating how to map portions of images subject to the second resolution allocation to a corresponding portion of the environmental model.

The methods and apparatus can be used for rendering stereoscopic images, e.g., pairs of images to be displayed to a users left and right eyes, or mono-scope images. Thus while the methods are well suited for use in simulating 3D environments they are also well suited for use in communicating panoramic images which may correspond to an area less than a full 360 degree environment and which may not be stereoscopic in nature.

Numerous additional methods and embodiments are described in the detailed description which follows.

While steps are shown in an exemplary order it should be appreciated that in many cases the order of the steps may be altered without adversely affecting operation. Accordingly, unless the exemplary order of steps is required for proper operation, the order of steps is to be considered exemplary and not limiting.

Some embodiments are directed a non-transitory computer readable medium embodying a set of software instructions, e.g., computer executable instructions, for controlling a computer or other device to encode and compresses stereoscopic video. Other embodiments are directed a computer readable medium embodying a set of software instructions, e.g., computer executable instructions, for controlling a computer or other device to decode and decompresses video on the player end. While encoding and compression are mentioned as possible separate operations, it should be appreciated that encoding may be used to perform compression and thus encoding may, in some include compression. Similarly, decoding may involve decompression.

The techniques of various embodiments may be implemented using software, hardware and/or a combination of software and hardware. Various embodiments are directed to apparatus, e.g., a image data processing system. Various embodiments are also directed to methods, e.g., a method of processing image data. In some embodiments, one or more of the method steps is implemented using a processor. Various embodiments are also directed to a non-transitory machine, e.g., computer, readable medium, e.g., ROM, RAM, CDs, hard discs, etc., which include machine readable instructions for controlling a machine to implement one or more steps of a method.

Various features of the present invention are implemented using modules. Such modules may, and in some embodiments are, implemented as software modules. In other embodiments the modules are implemented in hardware. In still other embodiments the modules are implemented using a combination of software and hardware. In some embodiments the modules are implemented as individual circuits with each module being implemented as a circuit for performing the function to which the module corresponds. A wide variety of embodiments are contemplated including some embodiments where different modules are implemented differently, e.g., some in hardware, some in soft-

45

ware, and some using a combination of hardware and software. It should also be noted that routines and/or sub-routines, or some of the steps performed by such routines, may be implemented in dedicated hardware as opposed to software executed on a general purpose processor. Such embodiments remain within the scope of the present invention. Many of the above described methods or method steps can be implemented using machine executable instructions, such as software, included in a machine readable medium such as a memory device, e.g., RAM, floppy disk, etc. to control a machine, e.g., general purpose computer with or without additional hardware, to implement all or portions of the above described methods. Accordingly, among other things, the present invention is directed to a machine-readable medium including machine executable instructions for causing a machine, e.g., processor and associated hardware, to perform one or more of the steps of the above-described method(s).

Numerous additional variations on the methods and apparatus of the various embodiments described above will be apparent to those skilled in the art in view of the above description. Such variations are to be considered within the scope.

What is claimed:

1. A method of communicating information to be used to represent an environment, the method comprising: operating a content delivery system to communicate to a content playback device, via a network interface, a first image map mapping portions of a first frame to segments of an environmental model, said first image map allocating different size portions of the first frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said environmental model; operating the content delivery system to communicate the first frame, including at least a portion of a first image to be mapped to said environmental model using said first image map, to the content playback device, said first frame corresponding to a first time; operating the content delivery system to communicate to the content playback device, via the network interface, a second image map mapping portions of a second frame to segments of the environmental model, said second image map allocating different size portions of the second frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said environmental model, said second image map allocating a different number of pixels to a first segment of said environmental model than are allocated by said first image map.

2. The method of claim 1, further comprising:

operating the content delivery system to communicate the second frame, including at least a portion of a second image to be mapped to said environmental model using said second image map, to the content playback device.

3. The method of claim 2, wherein said first frame and second frame correspond to different times; and

wherein a different portion of an environment to which the first and second frames correspond is important at the different times, said first and second image maps providing different resolution allocations to provide higher resolution to the important portion of the environment for each of said different times.

4. The method of claim 2, wherein said first and second frames include the same number of pixels.

5. The method of claim 2, wherein said first and second image maps are communicated in a content stream that includes said first and second frames.

46

6. The method of claim 2, further comprising: p1 operating the content delivery system to select a first resolution reduction operation to be performed on a first captured image to produce said first image; and

operating the content delivery system to select a second resolution reduction operation to be performed on a second captured image to produce said second image, said second resolution operation being different from said first resolution reduction operation.

7. The method of claim 6, wherein the first and second images correspond to the same portion of the environment.

8. A content delivery system including a network interface and a processor, said processor being configured to control the content delivery system to:

communicate to a content playback device, via a network interface, a first image map mapping portions of a first frame to segments of an environmental model, said first image map allocating different size portions of the first frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said environmental model; communicate the first frame, including at least a portion of a first image to be mapped to said environmental model using said first image map, to the content playback device, said first frame corresponding to a first time; and communicate to the content playback device, via the network interface, a second image map mapping portions of a second frame to segments of the environmental model, said second image map allocating different size portions of the second frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said model, said second image map allocating a different number of pixels to a first segment of said model than are allocated by said first image map.

9. The content delivery system of claim 8, wherein said processor is further configured to control the content delivery system to:

communicate the second frame including at least a portion of a second image to be mapped to said environmental model using said second image map to the content playback device.

10. The content delivery system of claim 9, wherein said first frame and second frame correspond to different times; and

wherein a different portion of an environment to which the first and second frames correspond is important at the different times, said first and second image maps providing different resolution allocations to provide higher resolution to the important portion of the environment for each of said different times.

11. The content delivery system of claim 9, wherein said first and second image maps map different numbers of pixels to an area corresponding to the same portion of an environment thereby providing different resolution allocations for said same portion of the environment based on which of the first and second image maps are used.

12. The content delivery system of claim 9, wherein the content delivery system includes a server providing a real time content stream while an event is ongoing.

13. A non-transitory computer readable medium comprising processor executable instructions, which when executed by a processor, control a content delivery system to: communicate to a content playback device, via a network interface, a first image map mapping portions of a first frame to segments of an environmental model, said first image map allocating different size portions of the first frame to differ-

ent segments of the environmental model thereby allocating different numbers of pixels to different segments of said model; communicate the first frame, including at least a portion of a first image to be mapped to said environmental model using said first image map, to the content playback 5 device, said first frame corresponding to a first time; and communicate to the content playback device, via the network interface, a second image map mapping portions of a second frame to segments of the environmental model, said second image map allocating different size portions of the 10 second frame to different segments of the environmental model thereby allocating different numbers of pixels to different segments of said environmental model, said second image map allocating a different number of pixels to a first segment of said environmental model than are allocated by 15 said first image map.

14. The non-transitory computer readable medium of claim 13, further including processor executable instructions, which when executed by a processor, control the content delivery system to: communicate the second frame 20 including at least a portion of a second image to be mapped to said environmental model using said second image map to the content playback device.

* * * * *